UNIVERSITY OF SOUTHAMPTON

FACULTY OF HUMANITIES

DEPARTMENT OF PHILOSOPHY

EPISTEMIC VALUE AND VIRTUE EPISTEMOLOGY

by

TSUNG-HSING HO

Thesis for the degree of Doctor of Philosophy

Month_2015

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF HUMANITIES

Philosophy

Thesis for the degree of Doctor of Philosophy

EPISTEMIC VALUE AND VIRTUE EPISTEMOLOGY

Tsung-Hsing Ho

My contributions to the research on epistemic value can be divided into two parts: first, I pinpoint some causes of the problems about epistemic value which have not previously been identified; and, second, I offer novel accounts of epistemic value which offer better solutions to the problems about epistemic value.

First, there are two trends in the literature on epistemic value that are rarely challenged: (i) epistemologists tend to understand epistemic value in terms of *intrinsic value from the epistemic point of view*, and (ii) the discussion of epistemic value tends to focus only on the values of properties of belief. I argue that both trends should be rejected if we want to solve several persistent problems about epistemic value: the value problems about knowledge, the teleological account of epistemic normativity, and the triviality objection that some true beliefs (or knowledge) are too trivial to be epistemic goals.

My account of epistemic value is in terms of *goodness of epistemic kinds*, which rejects (i). An epistemic kind is an evaluative kind—a kind that determines its own evaluative standards—whose evaluative standards are truth-directed: e.g. a belief is good *qua* belief if true. I argue that my account is immune from the triviality objection. Moreover, since the goodness of an epistemic kind is finally valuable, the account gives us simple solutions to the value problems of knowledge. I develop my own solutions through critically appropriating the virtue-theoretic account, according to which epistemic evaluation is a kind of performance evaluation, which rejects (ii). I argue that the value of knowledge consists of the value of epistemic success (true belief) and epistemic competence. Finally, I argue that approaches that focus on the evaluation of belief cannot explain epistemic normativity. Instead, we need an approach that focuses on the evaluation of *person*, which rejects (ii). I argue that conforming to epistemic norms is part of what makes us good *qua* person. The goodness of person *qua* person is an intrinsic value and able to provide *pro tanto* reasons for a person to be epistemically good *qua* person, which is the ground of epistemic normativity.

Overall, there are two main differences between my account and the mainstream account: first, the purpose of epistemic evaluation is about good cognitive performances rather than good beliefs; and, second, what grounds epistemic normativity is the goodness of a person *qua* person rather than the goodness (or correctness) of belief *qua* belief. The upshot of my account is that the focus of epistemology should be on questions such as 'What is an epistemically good person?' and 'What makes a person epistemically good *qua* person?' Furthermore, my account shows that epistemic normativity is not distinct from ethical normativity. That is, the question 'What is an epistemically good person?' is part of the question 'What is a good person?' and a reason why we should be an epistemically good person is consequently a reason why we should be a good person.

Table of Contents

List of Figures

(Permission to reproduce Figure 1 and 2 has been granted by Springer.)

DECLARATION OF AUTHORSHIP

I, Tsung-Hsing Ho,

declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

EPISTEMIC VALUE AND VIRTUE EPISTEMOLOGY

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. None of this work has been published before submission.

Signed: .....

Date:.....

Acknowledgements

Like every thesis, my thesis would be entirely different if I had a different supervisor. If every place where I received excellent advices from my supervisor is acknowledged, the thesis would be twice as long as it is (which I am not allowed to do that due to the length constraint). So I would like to express my sincerest gratitude to Daniel Whiting for his encouragement, patience, and support.

It is fortunate for me to come to Southampton when a growing number of philosophers come here doing excellent research on normativity. I am grateful to Conor McHugh, Jonathan Way, Lee Walters, Genia Schönbaumsfeld, Kurt Sylvan, Alex Gregory, and Ray Monk for their comments and suggestions. Also to my peers—Javier González, Sophie Keeling, James McGuiggan, Maria Mjaaland Sele, William Sharkey, Elizabeth Ventham, Krisen Jeffs—for our stimulating research seminars and organising the BPPA annual conference together. It would be less fun without them. I would also like to thank the external examiner, Anthony Booth. It is my loss that I know his works too late, with which my thesis shares a strong affinity.

For all these to happen, I must thank Tony Cheng for informing me the PhD studentship that I was awared. I am grateful to the University and the Humanities School for the studentship and to The Royal Institute of Philosophy for the bursary, without which the thesis could not be materialised. Also, I thank Springer to grant me permission to reproduce Figure 1 and 2 in Chapter 4.

When coming to Southampton, I bought a single one-way ticket and thought that I would return home only after my study is finished. I did not expect that I need to buy two more tickets now! I thank my wife, Qing, for her love and encouragement, and our son, Jiming, for all the laughter he brings to us. We would like to express our deepest gratitude to our dads and moms for their full support and assistance. Their visits make us feel that we are never away from home.

Introduction

Epistemic value has recently become a central topic in epistemology. In this introduction, I will first explain some of the reasons behind the resurgence of interest in epistemic values. I will then outline the contributions I hope to make to this topic.

The Value Turn in Epistemology

One of the fast growing literatures in epistemology is about epistemic value. This trend is even dubbed 'the value turn' in epistemology (Riggs 2007). One primary motive behind the value turn is the increasing dissatisfaction with the research on the Gettier Problem. For years after Edmund Gettier (1963) published his counterexamples to the then orthodox account of knowledge as justified true belief, the analysis of knowledge seemed to capture the full attention of epistemologists. With every new analysis of knowledge proposed, counterexamples soon followed. The whole cycle repeated itself again and again, leading to increasingly complex analyses and counterexamples. After several decades of failure of finding a satisfactory salutation, epistemologists started to reflect on the whole business of the Gettier problem. In Timothy Williamson's words, 'if the concept *knows* were the more or less ad hoc sprawl that analyses have had to become; why should we care so much about *that*?' (2000, 31).

Williamson intend to use his remarks as a modus ponens: if knowledge is valuable, the account of knowledge cannot be that ad hoc; intuitively, knowledge is valuable; hence, an analysis of knowledge cannot be ad hoc. The intuition that knowledge is valuable dates back to Plato. In Plato's *Meno*, Socrates asks Meno why knowledge is more valuable than mere true belief. Call it *the Meno problem*. Williamson's view suggests that an adequate theory of knowledge should

help us explain how knowledge is more valuable than mere true belief.[1] Beside the work of Williamson, another major approach heading in this direction is the virtue-theoretic account of knowledge: one knows that p, if and only if p is true and one believes that p because of one's intellectual abilities or virtues. Virtue epistemologists, like John Greco (2010), Ernest Sosa (2007; 2011), Wayne Riggs (2002; 2007) and Linda Zagzebski (1996; 2003), agree with Williamson that a theory of knowledge ought to explain why knowledge is valuable. They argue that the virtue-theoretic account of knowledge is better than other accounts, because it can provide satisfactory answers to the Meno problem, but others cannot.

Jonathan Kvanvig (2003)—who is also a virtue epistemologist—develops a more radical approach. Kvanvig interprets Williamson's remark, instead, as a modus tollens: if no account of knowledge can show why knowledge is valuable, then there is no need to care about solving the Gettier problem. Kvanvig does not deny that some theories of knowledge, such as virtue epistemology, can solve the Meno problem. But he argues that those theories of knowledge, at best, show only that some necessary elements of knowledge, besides true belief, are valuable, thus making knowledge more valuable than mere true belief. However, none of the theories can answer the problem of how knowledge is more valuable than any subset of its components (I call it *the Gettierian value problem* in the thesis). For example, Kvanvig argues that the virtue-theoretic account still encounters Gettier-style counterexamples, so it does not show how knowledge is more valuable than gettierised belief. Kvanvig's striking proposal is that no theory of knowledge can explain why knowledge is more valuable than any subset of knowledge's components.[2] Hence, Kvanvig argues that the focus in epistemology should not be on knowledge but on *understanding*, which Kvanvig regards as more valuable than knowledge. The nature and value of understanding has, since then, grabbed the attention of many epistemologists (Elgin 2006; Grimm 2012; Pritchard 2010).[3]

---

[1] Williamson's discussion of the Meno problem can be found in his (2000, 78-80). Jonathan Kvanvig (2003, 12-21) and Miranda Fricker (2009) critically discuss Williamson's view. Mark Kaplan's paper (1985) is probably the earliest one to bring the Meno problem back into contemporary epistemology.

[2] Ward Jones' paper (1997) is the forerunner of Kvanvig's argument

[3] Another interesting, but neglected, development is William Alston's epistemic desiderata approach (Alston 1993; 2005; see also Swinburne 2001). Alston argues that epistemologists should stop arguing which is the correct theory of epistemic justification, because each theory of justification is theorising different kinds of epistemic desiderata.

Another root of the value turn is the question concerning epistemic normativity. Discussion of epistemic normativity in contemporary epistemology is dominated by *epistemic teleology*.[4] Like its counterpart in ethics, epistemic teleology holds that the good is explanatorily prior to the right and attempts to explain epistemic normativity in terms of the intrinsic value of the consequences (usually the beliefs produced through cognitive procedures). The most popular account of epistemic teleology is *veritism* (Bonjour 1985; Foley 1987; Goldman 1999; 2002): the view that *truth* is the only epistemic goal or epistemic intrinsic value and *truth-conduciveness* is the only factor in determining whether one is epistemically justified or rational. Until recently, however, epistemic teleology was not subjected to the same close scrutiny as ethical teleology. The increasingly heated debate on epistemic teleology contributes significantly to the value turn in epistemology, since the theory of epistemic intrinsic value is the foundation of epistemic teleology.[5]

There are several objections to veritism, which will be looked into in the thesis. The first problem of veritism is that it is unable to answer the Meno problem. The reason is that the value of justification, according to veritism, is merely instrumental to the value of truth. Once a belief is true, justification is unable to add any value to the true belief. It is difficult, then, to see how knowledge is more valuable than mere true belief. Kvanvig (2003) calls this problem *the swamping problem* (Michael DePaul (2001) , Ward Jones (1997), and Linda Zagzebski (2003) offer similar arguments). That is, the value of justification or knowledge is swamped by the value of true belief, if veritism is true.

Second, veritism is often rejected on the grounds of *the triviality objection*: some true beliefs are about trivial matters, so they are not valuable at all. The triviality objection causes a problem for epistemic teleology to explain epistemic normativity: that is, since not all true beliefs are epistemically valuable, it cannot explain why we should, other things being equal, be epistemically justified in having *any* belief. The reason is that if some true beliefs are not epistemically valuable, it does not matter, from the teleologist perspective, that one is justified in having those beliefs because it is not conducive to any epistemic value (Kelly 2003; Grimm 2009). Other versions of epistemic teleology would also face the same objection. For example,

[4] Selim Berker (2013a) and Marian David (2001) show how widespread epistemic teleology is among epistemologists

[5] In the literature, the phrase 'epistemic final value' is used more often. As I will argue in Chapter 1, however, there should be a distinction between final value and intrinsic value, and what epistemic teleology needs is a theory of epistemic intrinsic value.

even if one holds that knowledge is the epistemic intrinsic value, it still faces the triviality objection (Baehr 2012).

The final problem is about epistemic teleology in general. Selim Berker (2013a; b) argues that epistemic teleology should be rejected, because, just as ethical teleology is thought to admit problematic trade-offs among person, epistemic teleology also admits problematic trade-offs among propositions. If believing a false proposition would lead me to many true beliefs—for example, if I believe in God, then the Church will allow me access to its secret archives, which makes me able to gain inside knowledge of the history of the Church—epistemic teleology seems to entail that my belief in God is justified. However, Berker argues that obviously my belief is not justified, so epistemic teleology should be rejected.

My Thesis

I have summarised two related trends in recent work on epistemic value: the value problems of knowledge and epistemic teleology. There are certainly other issues I have not discussed, but these two topics are the ones which have received the most attention in the literature and which receive critical treatment in this thesis. Although the research on epistemic value is booming, I argue in the thesis that there are several shortcomings in the research that have not yet been identified. As a result, I don't think that we have satisfactory solutions to the problems about epistemic value. And I believe that the solutions I propose in this thesis are both innovative and satisfactory. My contributions to this topic can be divided into two parts: first, I pinpoint some causes of the problems about epistemic value which have not previously been identified; and second, I offer novel accounts of epistemic value which offer better solutions to the problems about epistemic value. In the following, I will say more about each of these in turn.

The Causes of the Problems about Epistemic Value

There are two trends in the literature on epistemic value that are rarely challenged. I argue that those trends must be rejected if we want to solve the problems about epistemic value. Those trends are, first, that epistemologists tend to understand epistemic value in terms of *intrinsic value from the epistemic point of view*, and, second, that the discussion of epistemic value tends to be *belief-based*. Let me explain each in turn.

What is 'epistemic value'? Love and beauty are valuable and interest not only philosophers, but also ordinary people. However, they are not the kinds of values that particularly interest epistemogists. So epistemologists use the concept of epistemic value to distinguish the kind of value that interests epistemologists from other kinds of value, such as moral value or prudential. Interestingly, although the research on epistemic value is flourishing, epistemologists rarely give detailed accounts of the concept of epistemic value. When requested for some explanation, the most popular one is that epistemic value is, in William Alston's term (2005), an *intrinsic value from the epistemic point of view*.[6]

However, I argue in Chapter 1 that the account of epistemic value as intrinsic value from the epistemic point of view is the cause of the triviality objection. In Chapter 5, I argue that epistemologists who adopt this account of epistemic value will have difficulty in explaining why epistemic norms are categorical. To avoid those problems, the account of epistemic value as intrinsic value from the epistemic point of view should be rejected.

The second trend is as follows. The discussion about epistemic value tends to be belief-based. By that I mean that the focal point of epistemic evaluation is almost always *belief*. Take the Meno problem and the swamping problem for example. The problems are usually interpreted as 'Why is knowledge a more valuable kind of *belief* than mere true belief?' which naturally leads to the question '*Which property of belief* (except the property of being true) has epistemic value that would not be swamped by the value of truth?' In Chapter 3, I argue that the belief-based conception of epistemic evaluation misleads epistemologists to search for, in my term, *the Meno property of belief*: the property of belief that can increase the epistemic value of belief in addition to the value of truth. However, given the belief-based conception, the swamping problem is genuine, since there is no property of belief that can be plausibly said to increase the epistemic value of true belief.

It is widely held that virtue epistemologists, such as Greco and Sosa, can offer a satisfactory solution to the swamping problem. However, in Chapter 3, I show that, if their account remains belief-based, the virtue-theoretic account will also suffer from the swamping problem. Nevertheless, I think that their account is on the right track. In Chapter 2, I argue that the

[6] Michael Ridge (2013) is the only one who ever gives a detailed account of the concept of epistemic value. He says that epistemic value should be understood as a kind of attributive value rather than predicative value. He then suggests that epistemic value should be understood as something valuable

virtue-theoretic account of epistemic evaluation should not be taken as belief-based. In Chapter 3, I show how the virtue-theoretic account can offer straightforward solutions to the value problems of knowledge once the belief-based approach is rejected.

Furthermore, in Chapter 5, I argue that a belief-based account of epistemic evaluation is unable to explain epistemic normativity. That is, it cannot explain why epistemic norms, such as we should not have false or unjustified beliefs, are categorical, because there is no evaluative feature of belief that is normatively robust enough to generate categorical norms.

My Account of Epistemic Evaluation

Given the causes of the problems about epistemic value that I have identified, my account of epistemic evaluation is not belief-based and rejects the generally accepted view of epistemic value as intrinsic value from the epistemic point of view. The result is a novel account of epistemic value that offers straightforward and satisfactory solutions to the problems about epistemic value.

In Chapter 1, I propose my account of epistemic value in terms of *goodness of epistemic kinds*. An epistemic kind is an evaluative kind—a kind that determines its own evaluative standards—whose evaluative standards are truth-directed. For example, belief and perception are epistemic kinds. A belief is good *qua* belief if true; a perceptual faculty is good *qua* perceptual faculty if *reliably accurate*. I argue that my account is immune from the triviality objection. In Chapter 5, I develop an account of epistemic normativity based on my account of epistemic value.

Since the goodness of an epistemic kind is good for its own sake, the account gives us simple solutions to the value problems of knowledge. I develop the solutions through my critical appropriation of the virtue-theoretic account (Greco 2010; Sosa 2007; 2011), according to which epistemic evaluation is a kind of performance evaluation. According to Sosa, a belief is a kind of performance and can be evaluated as *successful* if true, *competent* if skillful or adroit, and *apt* if successful because competent. And when the belief is apt, it is knowledge. In Chapter 2, I defend the virtue-theoretic account of performance evaluation (that a performance can be evaluated in terms of success, competence, and aptness) in response to

---

from the epistemic point of view and what is of epistemic value is determined by what is the goal of inquiry.

the objection that performance evaluation is not applicable to belief, because belief is a state rather than a performance. I argue that performance evaluation is still applicable to belief even if belief is a state. The reason is that we often evaluate a performance through evaluating the end result of the performance, even if the performance was finished long ago. Belief is a product of a cognitive performance, so similarly we can evaluate the cognitive performance through evaluating the belief. When evaluating the belief in this way, we are not evaluating the belief *qua* belief (or mental state). Instead, what we are really evaluating is the cognitive performance (that produces the belief). Hence, one can accept the virtue-theoretic account, according to which epistemic evaluation is a kind of performance evaluation, and maintain the view that belief and knowledge are kinds of state.

The significance of my interpretation of the virtue-theoretic account becomes manifest in Chapter 3. My account is in opposition to the belief-based conception of epistemic evaluation. As I have said, the belief-based conception of epistemic evaluation holds that the value of knowledge consists entirely of the values of properties of belief. According to my account, however, the value of knowledge consists of the value of epistemic success (the goodness of a belief *qua* belief) and epistemic competence (the goodness of a cognitive ability *qua* cognitive ability). In Chapter 3, I show how the swamping problem is difficult to avoid if the belief-based conception is adopted and how my account offers a straightforward solution to the swamping problem.

In Chapter 4, I tackle the Gettierian value problem with my virtue-theoretic account of epistemic evaluation and argue that Greco and Sosa are wrong in holding that knowledge is more valuable than mere gettierised belief because apt (successful because competent) belief is more valuable than inapt, but successful and competent, belief. I argue that success and competence are the only two evaluative standards of performances, and aptness does not provide additional value to performances. Hence, I agree with Kvanvig and Pritchard that knowledge is no more valuable than gettierised belief. But their argument is based on the claim that the virtue-theoretic account still faces Gettier-style counterexamples. However, I argue that when the virtue-theoretic account is properly understood, it is able to cope with those counterexamples, but still its implication is that knowledge is no more valuable than gettierised belief.

I address an objection to my account of epistemic value in Chapter 5. One may object that the goodness of an epistemic kind is not normatively robust enough to explain epistemic

normativity. The fact that true beliefs are good qua belief does not give us a reason to have true beliefs. While it is obviously true, I argue that approaches that focus on the evaluation of belief cannot explain epistemic normativity as well. I argue that, to account for epistemic normativity, we need an approach that focuses on the evaluation of *person*. I argue that person is also an epistemic kind, that is, some of the evaluative standards of being good *qua* person are truth-directed. Based on my virtue-theoretic account, I suggest that a person is (epistemically) good *qua* person if she is epistemically competent or she has lots of true beliefs and few false beliefs (namely, the ratio of her true belief over false belief is high). Moreover, I argue that the goodness of a person is an intrinsic value and able to provide *pro tanto* reasons for us to be epistemically good *qua* person, which is the grounds of epistemic normativity. Hence, while my account of epistemic value cannot *directly* explain epistemic normativity, it provides the basis for my account of epistemic normativity.

Overall, there are two main differences between my account and the mainstream account. First, I argue that the purpose of epistemic evaluation is about good cognitive performances rather than good beliefs. Second, I argue what really grounds epistemic normativity is the goodness of person *qua* person rather than the goodness (or correctness) of belief *qua* belief. Each has some interesting implications. Since the purpose of epistemic evaluation is about good performances, I think that means that what epistemic evaluation assesses is whether a person is epistemically successful and competent when having a belief. That would mean that the focus of epistemology should be on questions such as 'What is an epistemically good person?' and 'How to make a person epistemically good *qua* person?' Furthermore, since what grounds epistemic normativity is the goodness of a person *qua* person, this means that epistemic normativity is not distinct from ethical normativity. Epistemic normativity is a subset of ethical normativity, which concerns the epistemic goodness of a person *qua* person. That is, the question 'What is an epistemically good person?' is part of the question 'What is a good person?' and a reason why should be an epistemically good person is consequently a reason why we should a good person.

One caveat. In the thesis, I assume that the virtue-theoretic account is correct. My aim is to show that it is not properly understood. This is important, because the virtue-theoretic account is the dominant epistemological view with regard to epistemic value. Moreover, I think that once the virtue-theoretic account is properly understood, its account of epistemic value is intuitively plausible and has interesting implications. However, one may reject my thesis if one rejects the virtue-theoretic account. That said, I do offer some responses to two

major objections to the virtue-theoretic account in Chapter 2 and the Appendix of Chapter 4.[7] So I believe that I have provided sufficient reasons why my account of epistemic value should be accepted.

Abstracts

Chapter 1. Epistemic Value and Epistemic Teleology

What is epistemic value? Although the literature concerning epistemic value is booming, the notion of epistemic value is rarely subjected to careful examination. Most epistemologists understand epistemic value as a kind of intrinsic value from the epistemic point of view. I call it 'the orthodox view' in this chapter. The orthodox view encourages *epistemic teleology*: an approach that attempts to explain epistemic normativity in terms of a theory of epistemic goal(s) (what is intrinsically valuable from the epistemic point of view) and a theory of epistemic justification (whatever is instrumentally valuable to achieve the epistemic goal(s)). A popular version of epistemic teleology is that truth is the epistemic goal and epistemic justification is understood in terms of truth-conduciveness. I argue that the orthodox view should be rejected, because it causes *the triviality objection* and encourages epistemic teleology, which suffers *the consequentialist problem*. The triviality objection is that some true beliefs are too trivial to be goals that are worth pursuing even from the epistemic point of view. The consequentialist problem is that epistemic teleologists have to consider *whatever* is instrumental to achieve the epistemic goal as epistemic justification. I explain how the orthodox view gives rise to these problems. I then propose my own account of epistemic value that avoids them. I account for epistemic value in terms of *goodness of epistemic kinds*. My account of epistemic value is based on a careful distinction between intrinsic value and final value, which is neglected by value theorists.

Chapter 2. Epistemic Evaluation as Performance Evaluation

This chapter is a preliminary to the next two chapters, in which I adopt the virtue-theoretic account to account for the epistemic value of knowledge. The virtue-theoretic account claims

[7] One major objection I do not discuss in the thesis is that the virtue-theoretic account is too strong (Lackey 2007; 2009; Vaesen 2011). For responses from virtue epistemologists, see (Carter forthcoming; Greco 2012; Kelp 2013).

that epistemic evaluation is a kind of performance evaluation, according to which evaluating beliefs is like evaluating a sport or musical performance. In this chapter, I critically examine the virtue-theoretic account through the objection that a belief cannot be evaluated as a kind of performance, because a belief is a state rather than a performance. On the one hand, I agree with the objection that, since belief is a state, performance evaluation, strictly speaking, is not applicable to belief. On the other hand, I argue that we often evaluate a performance through evaluating the result of the performance. Since a belief is a result of a cognitive performance, we can nevertheless evaluate a belief as a kind of cognitive performance when we want to evaluate the cognitive performance that produces the belief. Hence, performance evaluation is still applicable to belief. The upshot of my account is that when we evaluate a belief under performance evaluation, what we evaluate is not the value of belief *qua* belief, but the value of cognitive performance *qua* cognitive performance. Its implication on the value of knowledge can be seen in the next chapter.

Chapter 3. How Not to Swamp the Value of Knowledge

In the *Meno*, Socrates asked 'Why is knowledge more valuable than mere true belief?' Recently, Socrates' Meno problem has been reincarnated as the so-called *swamping problem*: 'How is the value of knowledge not swamped by the value of true belief?' The common diagnosis of the swamping problem is that the problem is caused by the assumption that truth is the only noninstrumental epistemic value. Call it *T-monism*. However, I argue that this diagnosis, while correct as far as it goes, does not go deep enough. By making a comparison between simple reliabilism and the virtue-theoretic account, I argue that the swamping problem is created by what I call *the belief-based conception of epistemic evaluation*: according to which, evaluating the epistemic value of knowledge is evaluating the values of the properties of belief. Once the belief-based conception is adopted, I argue that it is difficult to reject T-monism and the swamping problem looks insurmountable. However, based on the discussion in Chapter 2, I argue that epistemologists should adopt what I call *the performance-based conception*: according to which, evaluating the epistemic value of knowledge is evaluating the performance values of cognitive performance. When the cognitive performance results in knowledge, its values include the value of epistemic success (true belief) and the

value of epistemic competence, both of which are valuable for their own sakes. So the performance-based conception provides a straightforward solution to the swamping problem.

Chapter 4. Why Value Knowledge? Virtue Epistemology and the Gettierian Value Problem

Contrary to the claim of virtue epistemologists that the virtue-theoretic account can explain how knowledge is more valuable than gettierised beliefs (I call it *the Gettierian value Problem*), I develop a novel argument showing that, assuming that the virtue-theoretic account is right, it actually demonstrates that knowledge is no more valuable than gettierised beliefs. The virtue-theoretic account regards epistemic evaluation as a kind of performance evaluation: in Ernest Sosa's terms, a performance can be evaluated as *successful*, *competent*, and *apt* (*successful because competent*). The virtue-theoretic account considers knowledge as a kind of apt performance, and apt performances are more valuable than successful and competent (gettierised) performances. However, I argue that being apt does not make a performance better *qua* performance. The point of epistemic evaluation is about epistemic success and epistemic competence. The upshot is that, for the sake of epistemic evaluation, it is not necessary to distinguish between knowledge and gettierised belief, since both are equally successful and competent. An appendix is also added to address how the virtue-theoretic account could respond to objections from Gettier-style counterexamples based on some recent works by experimental philosophers.

Chapter 5. Epistemic Evaluation and Epistemic Normativity

In the final chapter, I try to sketch an account of epistemic normativity based on my account of epistemic evaluation. That is, I try to explain why some epistemic norms—such as the norm that we should not have false or unjustified beliefs—are categorical. I first criticise two prominent *evaluative approaches* to epistemic normativity: *the value approach* that attempts to explain epistemic normativity in terms of the epistemic value of belief, and *the correctness approach* that attempts to explain epistemic normativity in terms of the correctness of belief. I argue that both approaches cannot explain why epistemic norms are categorical. I then propose my account of epistemic normativity that attempts to explain epistemic normativity in terms of the goodness of person *qua* person. Call it *the personal value approach*. In a nutshell, I argue that (i) conforming to epistemic norms makes a person epistemically good *qua* person; (ii) being epistemically good *qua* person is part of what makes a person good *qua* person; (iii)

the goodness of a person *qua* person is intrinsically valuable, which provides a categorical reason to promote it; and therefore (iv) we should conform to epistemic norms. The upshot is that epistemic normativity is rooted in ethical normativity.

Chapter 1.

Epistemic Value and Epistemic Teleology

1.1 Introduction

The topic I want to pursue here is epistemic value. Although epistemic value has become one of the central topics in epistemology, there is not much discussion about how the notion of epistemic value should be understood. The discussion tends to proceed without a careful account of epistemic value.[8] For example, in their introduction to the edited book, *Epistemic Value* (Haddock et al. 2009), Adrian Haddock, Alan Millar, and Duncan Pritchard begin with this description of epistemologists' increasing interests in epistemic values,

> In recent years there has been a surge of interest in issues about *value in epistemology*. Two themes loom large. One is provided by puzzles about the value of knowledge. The starting point, but only the starting point, is provided by the question, 'How can knowledge be more valuable than merely true belief?' The other theme concerns *epistemic value* in a broad sense. Here the central issue is how best to make sense of epistemic appraisal conceived broadly to include evaluation of beliefs not just with respect to whether they amount to knowledge,

[8] A notable exception is Michael Ridge (2013). Ridge argues that epistemologists fail to heed the distinction between predicative and attributive goodness. Ridge's aim is to argue that if epistemologists want to use the epistemic value of knowledge as a pre-theoretical constraint on the theory of knowledge, epistemic value should be regarded as an attributive goodness. His reason that epistemic value is an attributive goodness is that 'being epistemically valuable (or good)' does not entail 'being valuable (or good). I think that Ridge is right. However, Ridge's point cannot settle the problems I discuss here. As I will explain in this chapter, a more crucial distinction is the one between final value and intrinsic value. The orthodox view holds that epistemic value is a kind of intrinsic value from the epistemic point of view. Since 'being valuable from the epistemic point of view' does not entail 'being valuable', the orthodox view, according to Ridge, also uses epistemic value as an attributive goodness. However, as we will see, the orthodox view is the cause of many problems in recent literature on epistemic value.

> but also with respect to whether they are, for instance, justified, or reliably formed, that is, formed through methods or processes that reliably yield true beliefs. (Haddock et al. 2009, 1; my italics)

Notice how swiftly the terminology is shifted from 'value in epistemology' to 'epistemic value'. 'Epistemic value' is a technical term that epistemologists use to define their topic as something different from moral value or prudential value. Most epistemologists understand the notion of epistemic value in terms of *a kind of value from the epistemic point of view* as opposed to moral or prudential value as a kind of value from the moral or prudential point of view. The phrase 'from the epistemic point of view' is coined by William Alston (1989; 2005). Call this account of epistemic value *the orthodox view*.

Talk of 'value in epistemology', in contrast, does not mark out a special type of value only from the epistemic point of view. It just refers to some values that, say, belief or knowledge has, which happen to be the subject matters of epistemology. That is, a value in epistemology could still be valuable, even if not from the epistemic point of view. Pritchard, in his later paper (2014), distinguishes between 'epistemic value' and 'the value of the epistemic'. 'The value of the epistemic' is actually another way to express 'value in epistemology'. The difference affects how we approach the question 'How can knowledge be more valuable than merely true belief?' An answer that knowledge is more practically valuable than merely true belief could be an answer to the question concerning the value of the epistemic, but not to the question concerning epistemic value.

Which kind of value should be the concern of epistemologists? Contemporary epistemologists primarily focus on epistemic values. Pritchard's arguments in that paper are basically that epistemologists' concerns are about epistemic value rather than the value of the epistemic. However, I think that epistemologists are wrong in thinking that they should be concerned with something only valuable from the epistemic point of view, rather than something epistemic that is just valuable. [9] One aim in this paper is to show that the orthodox view misleads epistemologists into thinking that there is a kind of value which is only valuable from

[9] For a good discussion on the value problem of knowledge both from the historical and contemporary perspectives without taking the problem as one concerning the epistemic value of knowledge, see John Hyman (2010). Linda Zagzebski (2003) discusses only the *value* of knowledge instead of the epistemic value of knowledge. Miranda Fricker (2009) also expresses doubt about the value problem being interpreted in a technical way.

the epistemic point of view. I will propose my own account of epistemic value, which effectively eliminates the distinction between epistemic value and value in epistemology.

Most of the discussion on the orthodox view focuses on how to understand the phrase 'from the epistemic point of view'. Alston understands the epistemic point of view in terms of *the epistemic purpose* (2005, 29). So the discussion focuses on what the epistemic purposes or goals are. The orthodox view, therefore, is a *teleological account of epistemic value*: an account that an epistemic value is either intrinsically valuable (that is, the epistemic goal) from the epistemic point of view or merely instrumentally valuable from the epistemic point of view. The orthodox view naturally encourages *the teleological account of epistemic justification* that accounts for epistemic justification purely in terms of epistemic instrumental value (henceforth, *epistemic teleology* for short).[10] To be clear, the orthodox view does not imply epistemic teleology. The orthodox view is concerning epistemic value, whereas epistemic teleology is a theory of epistemic justification. An epistemic teleologist must have a teleological account of epistemic value, but one can adopt the orthodox view without adopting epistemic teleology. Nevertheless, my observation is that epistemic teleology seems to be promoted by the orthodox view. Once one adopts the orthodox view, the only thing that makes one become an epistemic teleologist is to maintain that epistemic justification has only epistemic instrumental value. Epistemic teleology, however, faces two serious problems, which I will discuss in turn.

Most epistemic teleologists (Bonjour 1985; David 2014; Goldman 2002; Pritchard 2014) hold that truth is the one and only epistemic goal. Accordingly, true beliefs are the only things *intrinsically valuable* from the epistemic point of view and epistemic justification is thus understood entirely in terms of truth-conduciveness. Beliefs that are produced through whatever truth-conducive way would be considered epistemically justified, according to epistemic teleology.

The most common objection to epistemic teleology is *the triviality objection* that not all true beliefs are valuable; apparently, trivial true beliefs are not valuable at all and should not be our epistemic goal.[11] The triviality objection leads many epistemologists to look for other candidates for the epistemic goals. For example, Timothy Williamson (2000) and Declan

[10] Selim Berker (2013a; b) offers a more detailed analysis of the structure of epistemic teleology.
[11] The triviality objection is now so common that Jeremy Fantl and Matthew McGrath (2009, 166) call it the stock objection to the value of truth.

Smithies (2012) hold that the epistemic goal is knowledge. Jonathan Kvanvig (2003; 2014) is a pluralist about epistemic goals: truth, knowledge, understanding, etc. However, it is unclear how these approaches can meet the triviality objection to knowledge, understanding, etc.: if a true belief about trivial matters is not epistemically valuable, why is the knowledge or understanding of them epistemically valuable?[12] More importantly, as I will argue later, my point is not that the orthodox view that the triviality objection cannot be met, but rather that because of the orthodox view, the whole debate about the triviality objection can only end up in a deadlock. For the sake of simplicity, therefore, I henceforth assume that truth is the epistemic goal for epistemic teleology.

We have seen how the orthodox view encourages epistemic teleology and turns questions about epistemic value into questions about epistemic goals. However, even if the triviality objection can be solved, there is a more radical objection by Selim Berker (2013a; b): that is, epistemic teleology is untenable, because epistemic teleology admits problematic trade-offs. Just like moral teleology may imply that one can justifiably kill one innocent person in order to save several people, epistemic teleology may also imply that one can justifiably have an unfounded false belief because having that false belief leads one to have many true beliefs. Call it *the consequentialist problem*.[13]

One aim of this chapter is to show that the orthodox view is the root of the triviality objection and the consequentialist problem about epistemic teleology. And those problems seem unable to be fixed by answering the question as to what the epistemic goal is. I would like to suggest that epistemologists are looking for the answer in the wrong place. The real problem about the orthodox view is not what the epistemic point of view or the epistemic goal is, but that it

[12] Jason Baehr (2012) and Daniel Whiting (2012b) run different versions of the triviality objection to knowledge and understanding. Some epistemologists hold views similar to Alston that the epistemic goal is 'to acquire true rather than false beliefs about matters that are of interest or importance to us' (2005, 29). Alston's view can overcome the triviality objection, but it faces serious problems (Grimm 2008; 2009; Kelly 2003) about how to make sense of epistemic normativity, such as why we should not have false beliefs, even if those beliefs are about trivial matters. I discuss more about this problem in the fifth chapter.

[13] Pluralists like Kvanvig may also regard justification as an epistemic intrinsic value. However, since justification belongs to the theory of epistemic justification, it seems wrong to include justification within the theory of epistemic value. Moreover, once the theory of epistemic justification is specified, such as reliabilism, the same problem arises again: an unreliable belief would also be considered justified if it leads to many reliable beliefs. There are defences of epistemic teleology (Ahlstrom-Vij and Dunn 2014; Goldman 2015) against Berker's criticisms. However, their defences mostly address Berker's counterexamples rather than the structural problem of epistemic teleology on which Berker focuses. In section 5, I will identify what the structural problem is.

regards epistemic value as intrinsic value (though merely from the epistemic point of view). I will explain in the following sections (section 3 and 5) how the view of epistemic value as intrinsic value promotes epistemic teleology, which suffers the triviality objection and the consequentialist problem.

Moreover, I will propose my own account of epistemic value (in section 4). An adequate theory of epistemic value, I think, should satisfy the following desiderata:

(D1) It should explain why the value(s) is *epistemic* rather than non-epistemic;

(D2) It should tell us what is epistemically valuable *for its own sake*;

(D1) should be a truism; after all, it is a theory *of epistemic value* rather than something else. (D2) may need more explanation. Philosophers are usually more interested in values for their own sakes rather than for the sake of other things. One reason is theoretical, because what is of value only for the sake of something else derives its value from something else. The things which are of value for their own sake are regarded as foundational and more important. Epistemologists are no different. The orthodox view can certainly satisfy (D1). But (D2) is where the orthodox view begins to go astray, because while intrinsic value is valuable for its own sake, not all things which are valuable for their own sakes are intrinsically valuable. That is, intrinsic value is a species of value for its own sake. The orthodox view tries to satisfy (D2) by holding a stronger theory of epistemic value than an adequate theory of epistemic value needs to be, which is the cause of the triviality objection and the consequentialist problem, as I will argue in this chapter.

The confusion of epistemologists between intrinsic value and value for its own sake is understandable, because this distinction is not properly understood even among philosophers in general. My claim would surprise many philosophers who are familiar with Christine Korsgaard's (1983) distinction between value in virtue of its own properties and the value for its own sake (in her term, final value), since her distinction is well accepted not only by non-epistemologists (Kagan 1998; Rabinowicz and Rønnow-Rasmussen 2000), but also by epistemologists (Brogaard 2006; Pritchard 2007). However, as I will argue in section 2, Korsgaard's distinction actually obscures the conception of intrinsic value. To clear up the confusion, I need to take a detour in the next section into value theory. Readers should heed that my distinction between intrinsic and final value is different from Korsgaard's distinction and the difference is not merely verbal.

1.2 On the Distinction between Intrinsic Value and Final Value

In this section, I want to clarify the kind of value philosophers used to call *intrinsic value* (e.g. Moore 1993a). Intrinsic value used to be (and perhaps still is) central to ethical theories. The idea of intrinsic value is something philosophers have long been familiar with, so it may seem strange there is any need for more clarification. Especially after Korsgaard's well-accepted distinction between value in virtue of its intrinsic properties and value for its own sake, philosophers now accept that value for its own sake—rather than value in virtue of its intrinsic properties—is the kind of value philosophers used to call it intrinsic value.[14] My aim in this section is to show that this received view is wrong, because not every value for its own sake is intrinsically valuable. Intrinsic value is a species of value for its own sake

Let me begin with how philosophers used to think about intrinsic value. G. E. Moore claims that an investigation of intrinsic value is 'an enquiry which belongs only to Ethics' (1993b, 57). W. D. Ross holds that things of such value fall within the 'scope of ethics': they are 'worthy objects of admiration' and 'worthy objects of satisfaction' (1939, 290). Roderick Chisholm (1978; 1981) analyses intrinsic value in terms of 'ethical requirement' and 'intrinsically preferability': roughly, p is intrinsically better than q, if and only if there is an ethical requirement for one to prefer p over q. Noah Lemos holds that whether we have an appropriate emotional attitude toward things of such value is an ethical matter and explains that value in terms of 'being intrinsically worthy of love' (1994, 12). In sum, Michael Zimmerman (1999) suggests that we call it *ethical value*. Throughout this paper, however, I will still call it intrinsic value. The point is that intrinsic value is something philosophers consider of great ethical importance.

Historically, philosophers use several terms to refer to intrinsic value, such as 'final value', 'value in itself', 'value for its own sake', or 'value as an end', and contrast it with 'extrinsic value', 'instrumental value', or 'value or the other's sake'. Korsgaard (1983), however, rightly argues that these terms convey different concepts of value. She argues that those terms express two pairs of values: first, final value (value for its own sake or as an end) versus

[14] Korsgaard, not without good reasons, chooses to call value in virtue of its intrinsic properties 'intrinsic value'. However, as we will immediately see, value in virtue of its intrinsic properties is not the kind of value I call 'intrinsic value'. To avoid confusion, I choose to use the wordy phrase 'value in virtue of its intrinsic properties' for Korsgaard's 'intrinsic value' in order to distinguish it from the kind of intrinsic value I want to clarify.

instrumental value (value for others' sakes or as a means); second, value in virtue of its own intrinsic properties versus value in virtue of its extrinsic properties. Philosophers before Korsgaard usually believed it to be a distinction without difference, because things valuable for their own sakes are valuable in virtue of their intrinsic properties and things valuable for others' sakes are valuable for their extrinsic properties. But Korsgaard argues that things could be finally valuable in virtue of their extrinsic properties, which is supported by several philosophers through the following examples: mink coats finally valuable for their luxury (Korsgaard 1983), a well-designed racing car or fine culinary skills finally valuable for their causal, instrumental properties (Kagan 1998), and Princess Diana's dress finally valuable for its causal history to a beloved woman (Rabinowicz and Rønnow-Rasmussen 2000). It may remain disputed whether there could be things finally valuable in virtue of their extrinsic properties,[15] but, as a *conceptual* distinction, it seems universally accepted. It is the conceptual distinction that I want to discuss here, though, for the present purpose, I will assume that those examples are finally valuable. The point is that these examples are used by philosophers to illustrate Korsgaard's notion of final value.

Korsgaard's distinction aims to clarify philosophers' talks of intrinsic value, the kind of value central to ethics, in Korsggard's term, 'that marks out as worthy of choice' and 'meaning that here and now the world is a better place because of this thing' (1983, 169). After Korsgaard made the distinction, some philosophers have held that the notion of value in virtue of its intrinsic properties 'may not be as normatively interesting as many have thought' (Rabinowicz and Rønnow-Rasmussen 2000, 48) and that 'it is with the notion of something's being valuable for its own sake [final value] that philosophers have traditionally been, and continue to be, primarily concerned' (Rønnow-Rasmussen and Zimmerman 2005, xxiv). In other words, it is *final value rather than value in virtue of its intrinsic properties* that is the intrinsic value that plays a central role in ethics.

This received view is not entirely right, however. True, value in virtue of its own properties is not intrinsic value. But final value *per se*, I will argue, is not intrinsic value, either. Intrinsic value, I will argue later, is *a species of final value*, so that not every final value is intrinsically valuable. Korsgaard's distinction, therefore, fails to capture the notion of intrinsic value that it aims to clarify. A further distinction in final value is required if we do not want to mistake

[15] Michael J. Zimmerman (2001) denies the possibility that things could be finally valuable in virtue of their extrinsic properties. However, he accepts the distinction.

everything of final value as intrinsically valuable. The distinction is crucial to the present purpose, because the desiderata (D2) of a theory of epistemic value only requires it to tell us what is epistemically finally valuable rather than what is epistemically intrinsically valuable. If intrinsic value is a species of final value, the orthodox view is holding a needlessly strong account of epistemic value. Since Korsgaard's distinction obscures the notion of intrinsic value rather than clarifies it, it is not surprising that epistemologists are not aware of this problem.

Now I should explain why intrinsic value is just a species of final value. Let us first look at a selection of intrinsic values from William Frankena (1973): life, health, pleasure, knowledge, wisdom, beauty, virtues, friendship, freedom, peace. Compare this list with the above list of final value: a luxurious mink coat, a well-designed racing car, a fine culinary skill, and Princess Diana's dress. It is difficult to understand why an ethicist—since intrinsic value is, in Zimmerman's terms, ethical value—would put those things in the list of final value into a list of ethical value.[16] It is unlikely to be the case that mink coats and good culinary skills are the primary concerns of ethicists. The sharp contrast between the lists should make us reconsider whether the notion of final value is really the notion of intrinsic value with which ethicists have traditionally been concerned.

To see the difference between final value and intrinsic value, it is helpful to see how philosophers determine whether something is intrinsically valuable. Many philosophers (Bradley 2006; Brentano 1969; Broad 1930; Chisholm 1978; 1981; Ewing 1948; Moore 1993b; Ross 1939) try to detect things of intrinsic value by running two isolation tests (recall that Korsgaard also uses similar ideas to explain the kind of value her distinction aims to clarify):

> ONTOLOGICAL: Adding something of intrinsic value to the world, *ceteris paribus*, makes the world better;
>
> INTENTIONAL: If one evaluates something as such, without reference to circumstances or consequences, and evaluates it as intrinsically valuable, one must hold that there is a *pro tanto* reason for us to promote it, preserve it, or bring it about.

[16] There is no doubt that Korsgaard and Kagan are ethicists and they consider mink coats and culinary skills as finally valuable. And I agree that they are finally valuable. My point is that they are not ethically or intrinsically valuable (note that I do not use 'intrinsic value' in the way Korsgaard uses it), so that the notion of final value is equivalent to that of intrinsic value.

G. E. Moore (1993b, 133-36), for example, (in)famously holds that beauty passes both tests: beauty is intrinsically valuable because, *ceteris paribus*, adding beauty in the world would make the world better, and everyone, when contemplating beauty for its own sake, would agree that we all have a reason to preserve the beauty in the world or to make the world more beautiful. Some philosophers (Bodanszky and Conee 1981; Lemos 1994) have some doubts about the usefulness of the isolation tests. For the present purpose, however, we do not need to worry about that. All we need from the tests is to help us grasp the notion of intrinsic value those philosophers have in mind. Plausibly, things that pass the isolation tests are ethically valuable. Since things of intrinsic value make the world better and provide us reasons to promote or preserve it, it is understandable why such values lie at the heart of ethics.

But can final values pass the isolation tests? It seems unlikely. While it seems plausible that, *ceteris paribus*, the more pleasure or wisdom there is, the better the world is, it seems implausible that a world where Princess Diana's dress or a Formula One car exists is better than a world where they do not. Similarly, although it seems right that, *ceteris paribus*, everyone has a reason to make the world more peaceful and beautiful, it seems false that everyone has a reason to produce or preserve as many mink coats as possible. Still, it sounds right that, *ceteris paribus*, everyone should be virtuous, but wrong that everyone should become skilled in cooking. If so, while intrinsic values are clearly valuable for their own sake, not every final value is intrinsically valuable. Hence, intrinsic value is a species of final value.

A warning: claims about intrinsic value are always disputable. One may have different judgments on the results of the tests. However, it should be manifest that some examples of final value cannot be plausibly said to pass the isolation tests. If that is the case, it is evidence that *some* final values are not intrinsically valuable. Indeed, even Woldek Rabinowicz and Toni Rønnow-Rasmussen, who consider Princess Diana's dress finally valuable, have doubts about whether it can pass the tests:

> We take Diana's dress to be [finally] valuable, but do we think it is valuable that this dress exists? We might hold that, because of its value, the dress should be preserved or that it should not be destroyed, but do we think that the world is a better place because of the existence of that object? (Rabinowicz and Rønnow-Rasmussen 2000, 43)

Clearly, Rabinowicz and Rønnow-Rasmussen doubt that Princess Diana's dress can pass ONTOLOGICAL, but they appear to think that it can pass INTENTIONAL. Nevertheless, they

merely claim that they *might* hold that. By contrast, Moore has no hesitation in claiming that beauty passes both tests. The hesitation of Rabinowicz and Rønnow-Rasmussen about whether Princess Diana's dress can pass the tests is evidence that they do not consider the final value of the dress as intrinsically valuable, since intrinsic value is meant to be able to pass the isolation tests.

Furthermore, as Rabinowicz and Rønnow-Rasmussen recognise, if something can become finally valuable because of its special relation to a particular person, object, or event, then there will be many instances of final value. If final value *is* intrinsic value, then it follows that we all have *pro tanto* reasons to promote or preserve all things of final value. But that seems absurd. Suppose that my life is valuable for its own sake, my birth, given its causal history to my life, may also be judged as finally valuable. Do we have reasons to preserve all things that have a causal history to my birth? Surely 'No'! It shows that not all final values are intrinsically valuable. In other words, final values can come cheap. It is hard to see how things can become intrinsically valuable so easily.

To be clear, it does not mean that things of final value cannot provide reasons. Plausibly, the fact that the dress belonged to Princess Diana could provide reasons for *some* people—perhaps for those who love and respect Princess Diana—to preserve her dress. And things related to my birth may provide reasons for my family and me to preserve them. However, INTENTIONAL requires that things of intrinsic value can provide reasons for *everyone* to promote them. If Princess Diana's dress cannot provide everyone reasons to preserve it, its final value is not an intrinsic value.

Therefore, not every final value is an intrinsic value. And since all intrinsic values are valuable for their own sake, all intrinsic values are finally valuable. Hence, intrinsic value is a species of final value. It is wrong to identify final value with intrinsic value.

So far my argument is based on intuition. Now I will provide a more principled argument for the thesis that intrinsic value is a species of final value. Let me first consider an objection: one may accept my judgment that Princess Diana's dress or a Formula One car is not intrinsically valuable (in the ethical sense), but arrive at a different conclusion. What my intuitions about the isolation tests instead show, one may think, is that those items are not valuable for their own sake, after all. Hence, there is no need to distinguish intrinsic value from final value. To this view, my response is that philosophers who follow Korsgaard are right to consider them finally valuable. The reason is that, although they could not pass the isolation tests, they could

offer everyone reasons to *appreciate* (but not *promote*) them. Knowing that the dress was worn by Princess Diana or the car is a superb racing car could certainly offer us reasons to have some positive responses toward them, but it does not thereby follow that we have reasons to promote or preserve them. In other words, final values can pass another form of the isolation test:

> INTENTIONAL*: If one evaluates something as such, without reference to circumstances or consequences, and evaluates it as finally valuable, one must hold that there is a *pro tanto* reason for us to respond to it positively.

Now, it is clear that intrinsic value is a species of final value, because promotion is one kind of positive response. A Formula One car can give us a reason to appreciate it as a good racing car, but not to preserve it or to produce more Formula One cars. Even something that is intrinsically bad can give us a reason to positively respond to it. For example, an atomic bomb is intrinsically bad, but one can still appreciate it as a powerful weapon. Things of final value only need to elicit some positive responses, but things of intrinsic value require being promoted. Therefore, something could be finally valuable, but not intrinsically valuable, though something intrinsically valuable must be finally valuable.

We can put the above point in terms of the popular *fitting attitude account of value* (INTENTIONAL and INTENTIONAL* are actually its non-reductive versions). According to the fitting attitude account, to be valuable, roughly, is to be a fitting object of some positive evaluative attitudes. Thus, we can draw the distinction between final value and intrinsic value based on different kinds of positive attitude. Final values are fitting objects of positive evaluative attitudes *in general*, while intrinsic values are fitting objects of some specific types of positive attitudes (e.g. consideration of worthiness of promotion).[17] This also highlights why

[17] Michael J. Zimmerman, in his entry for intrinsic value (2014) in *The Stanford Encyclopedia of Philosophy*, has an interesting discussion of the fitting attitude account of intrinsic value: 'It would thus seem very natural to suppose that for something to be intrinsically good is simply for it to be such that it is fitting to value it for its own sake. ("Fitting" here is often understood to signify a particular kind of moral fittingness, in keeping with the idea that intrinsic value is a particular kind of moral value.)' Although he initially says that to be intrinsically valuable is to be fitting to value for its own sake, he immediately emphasises that the fittingness in question is a kind of moral fittingness. There are certainly non-moral types of fittingness. So we can still distinguish between intrinsic value as moral fittingness to value for its own sake and final value as fittingness in general to value for its own sake. However, Zimmerman seems unaware the distinction between final value and intrinsic value. He remains thinking that Korsgaard's 'final value' is intrinsic value, as he says: 'Presumably ['intrinsic value'] has been used by many writers simply to refer to what Korsgaard calls final value, in which case the moral significance of (what is thus called) intrinsic value has of course not been thrown into doubt.'

intrinsic value is ethically valuable, because intrinsic values require us to take some action, that is, bring intrinsic values into the world or to prevent them from destruction. But final values cannot necessarily impose such requirements on us.

I have argued that the distinction between intrinsic value and final value has not been clearly drawn by philosophers. Indeed, Korsgaard's well-received distinction may even obscure the notion of intrinsic value. The confusion is understandable, because intrinsic value is also finally valuable. Therefore, we need to be careful to discern which value epistemologists are talking about when they claim that something is epistemically valuable for its own sake. The claim could mean that something, from the epistemic point of view, is finally valuable or intrinsically valuable. However, in the next section, I will argue that theorising epistemic value in terms of intrinsic value is problematic in several ways.

1.3 Truth as Intrinsic Value: the Orthodox View and the Triviality Objection

In this section, I aim to explain why we should not adopt the orthodox view on epistemic value. To show why this is the case, I will show that the triviality objection arises only because epistemic value is understood as a kind of intrinsic value, and the debate on the triviality objection would be inevitably locked in stalemate.

Let me begin with a view that every belief, simply and only because of being true, is intrinsically good. Call it *veritism*. A standard objection in the literature to veritism, mentioned above, is *the triviality objection*: some true beliefs are so trivial that they appear worthless. Beliefs concerning the truth about who are those people now blinking in the world or about the length of the grasses of the lawn seem hardly worthwhile. Perhaps it is not bad to believe them, but they seem to have no value at all.

To respond to the triviality objection, veritists usually respond that true beliefs are *prima facie* good, not all things considered good (Fantl and McGrath 2009, 165-67; Horwich 2006; Kvanvig 2003, 41; Kvanvig 2008, 201; Lynch 2005, 46-57; 2009a; b). Veritists admit that trivial truths are usually not worthy of being believed. However, that does not show that trivial true beliefs are not intrinsically valuable. The reason why we regard them as unworthy is that when we consider whether to have those beliefs, we need to weigh their value with the cost of having them. And because the value of trivial true beliefs is so little that it is easily overridden by the

cost of having them, we are thus misled to think that trivial true beliefs have no intrinsic value, though actually they do have some, however little. As Jonathan Kvanvig says:

> We are finite beings, with limited time and resources for enhancing our well-being; and without some special situation in which counting the grains bring pleasure to a person, perhaps only by passing the time in a way not completely boring, our general interest in enhancing our well-being comes into conflict with our general interest in the truth. Perhaps it is even true that most of the time when our interest in enhancing our well-being conflicts with our interest in the truth, the former overrides the latter. In any case, there is no obstacle to interpreting [the triviality objection] in this way, and if we do, we leave intact the earlier point, that obtaining the truth is valuable in itself, apart from any contribution it makes to our well-being. (Kvanvig 2003, 41)

However, the response is begging the question.[18] To see this point, run the tests ONTOLOGICAL and INTENTIONAL on trivial true beliefs, since the question is about whether all true beliefs are intrinsically valuable. The triviality objection can be interpreted as arguing that trivial true beliefs cannot pass the isolation tests. It seems true that a world does not become better simply because someone believes every truth about all the characters and plots in *EastEnders*. It also seems true that not everyone has a reason to make anyone believe any truth about *EastEnders*. Veritists want to argue that trivial true beliefs have little intrinsic value, but are often overridden. However, in both tests, the beliefs are evaluated in isolation; nothing can override the intrinsic value, if any, of trivial true beliefs. So this line of argument cannot persuade people who think that trivial true beliefs have no intrinsic value, because they cannot pass the isolation tests.

This objection to veritism is certainly not decisive. Veritists can insist that trivial true beliefs can nevertheless pass the tests. However, I think that this response is hardly convincing. It is in this dialectical context, I think, that veritists introduce the phrase 'from the epistemic point of view'. So veritists can respond to the triviality objection that, while trivial true beliefs are not intrinsically valuable *simpliciter*, they are intrinsically valuable *merely from the epistemic point of view*—thus we have the orthodox view. This is how Pritchard (2014) responds to the triviality objection. Pritchard, as mentioned above, makes the distinction between *epistemic*

[18] Jason Baehr (2012) and Daniel Whiting (2013) give different arguments why this response does not work.

value and *the value of the epistemic*: an epistemic value is something valuable only from the epistemic point of view, and the value of the epistemic concerns whether the epistemic value is valuable *simpliciter* or from other points of view. Hence, the orthodox view is holding a weaker claim: while trivial truths are not intrinsically valuable *simpliciter*, trivial truths are intrinsically valuable *from the epistemic point of view*.

The phrase 'from the epistemic point of view', therefore, is used not only as a way to confine the issue to a merely epistemological one, but also as a hedge to defend the evaluative claims. But this move makes the debate murkier than it already is. For the question now becomes whether trivial true beliefs makes the world better *from the epistemic point of view* or whether trivial true beliefs by themselves provide us reasons to promote them *from the epistemic point of view*. Since the phrase 'from the epistemic point of view' is purely a technical invention,[19] it is difficult to say whether our intuition is reliable enough to answer these questions. Indeed, several philosophers would reject the view that all true beliefs are intrinsically valuable even merely from the epistemic point of view. For example, Alston, who coined the phrase 'from the epistemic point of view', holds that the epistemic purpose is 'to acquire true rather than false beliefs about matters that *are of interest or importance to us*' (Alston 2005, 29; my italics). For Alston, trivial true beliefs are not valuable at all even from the epistemic point of view. Again, the debate ends up in a standoff.

Proponents of the orthodox view could respond, like veritists do, that even from the epistemic point of view, the intrinsic value of trivial true beliefs is *negligible*, so we might be misled into thinking that they are not valuable at all. But this move is particularly problematic, because the whole debate boils down to these questions: 'Does a trivial true belief make the world *slightly better*, and does a trivial truth give us *slight* reason to think that we should believe it?' Arguably, we do not have reliable intuitions to answer them. Not to mention that we also need to answer these questions *from the epistemic point of view*. As Jason Baehr comments, this response is 'blatantly question-begging. ... The problem, however, is that if the [triviality] objection cannot be defended in this way, it is far from clear how else it might be supported' (2012, 6).

[19] Many philosophers (e.g. Ahlstrom-Vij and Grimm 2013; Pritchard 2014) explain 'the epistemic point of view' in terms of 'the goal of inquiry', that is, what is valuable from the epistemic point view is what satisfies the goal of inquiry. However, the value of satisfying the goal of inquiry needs not be intrinsic. As I explain in note.14, assuming that truth is the goal of inquiry, an inquiry that attains truth is good *qua* inquiry, which is merely finally valuable.

In this section, I explain how the orthodox view that epistemic value is a kind of intrinsic value inevitably ends up in clashes between intuitions. This is a well-known problem about intrinsic value. As Kvanvig says, 'I grant that claims of intrinsic value are deeply troublesome claims to defend, but one cannot philosophise about normative areas of philosophy without such claims, and I see no way to defend a claim of intrinsic value other than in the manner done here' (2003, 42). I do not argue that we should give up talking about epistemic intrinsic value just because intuition clashes inevitably arise. My reason is that the talk of epistemic intrinsic value requires us to judge whether something has the slightest amount of intrinsic value and to judge it from the epistemic point of view. This makes the whole debate more problematic than the usual debates about intrinsic value in general. If the debate is nothing but a question about whether truth is at least slightly valuable, one may wonder whether it is worth any debate.

I have explained how the triviality objection inevitably arises when the orthodox view understands epistemic value as a kind of intrinsic value. But I have not explained why the consequentialist problem is insurmountable for the orthodox view. Before that, however, in the next section, I will propose my own account of epistemic value. I account for epistemic value as a kind of final value, which would be immune to the triviality objection.

1.4 Epistemic Value as Goodness of Epistemic Kind

Epistemologists are interested in what is epistemically valuable for its own sake, namely, what is epistemically finally valuable. As I have shown, intrinsic value is a species of final value. A theory of epistemic value needs not to be a theory about intrinsic values. In this section, I will propose an account of epistemic value as a kind of final value, but not intrinsic value.

My account of epistemic value is to understand epistemic value as *goodness of its own kind*: a true belief is good *qua* belief, a keen eyesight is good *qua* eyesight (cognitive ability), an intellectual virtue is good *qua* cognitive trait, etc. In the following, I will explain why goodness of its own kind is finally valuable (though not necessarily intrinsically valuable) and how my account is an account of *epistemic* value.

The idea of goodness of a kind (Foot 1978; 2001; Geach 1956; Kraut 2007; Thomson 2001; 2008; Williams 1972) is familiar. Many things can be evaluated in terms of goodness of their own kinds. There are kinds that themselves set the goodness standards that their instances have to meet so as to be good *qua* instances of their own kinds. Call such kinds *evaluative kinds*. Paradigmatic examples of evaluative kinds are functional kinds. The function of knife is

to cut, so a knife sharp enough to cut things easily is good *qua* knife. I remain open as to whether there are non-functional evaluative kinds.

Is goodness of a kind finally valuable? It is, because it can pass INTENIONAL*. For example, since the function of a knife is to cut, one evaluative standard of a knife is sharpness: namely, a sharp knife is good *qua* knife. Let us run the test INTENTIONAL* on a sharp knife. Presumably, the fact that the knife is good *qua* knife does not give us reasons to want it, to preserve it (to maintain its sharpness), or to produce more sharp knives. But when we recognise that being a sharp knife is good *qua* knife, it gives us reasons to *appreciate* it as a good knife. Appreciation is certainly a kind of positive response. The goodness *qua* knife thus passes INTENTIONAL*, and hence is finally valuable. The general lesson is that the goodness of a kind is a kind of final value, because it gives us reasons to appreciate their instances when they are good *qua* their own kinds.

True belief is finally valuable is the same way. Belief is also an evaluative kind. One of its evaluative standards is *truth*: a belief is good *qua* belief if it is true.[20] I think that this claim is plausible, and I will not provide any argument for it. The aim of this paper is not to articulate what is the evaluative standard of belief, but rather to show that my account of epistemic value is better than the orthodox view.

One advantage of my account over the orthodox view is that my account is immune to the triviality objection. If one wants to raise the triviality objection to my account, the objection would become that trivial true beliefs are not good *qua* beliefs. But this is obviously wrong. However trivial the beliefs are, we should still appreciate them as good instances of belief. The value of true belief *qua* belief is thus finally valuable. My account of epistemic value as goodness of its own kind, therefore, satisfies the desiderata (D2): it explains how the epistemic value of true belief is valuable for its own sake, and how trivial true beliefs are still epistemically valuable.

How about (D1)? Even though the value of true belief *qua* belief is finally valuable, I have not explained why it is an *epistemic* value. We should bear in mind that the notion of epistemic value is purely technical. Epistemologists use this term in order to isolate a kind of value that interests them. The phrase 'from the epistemic point of view' is meant to serve this purpose.

[20] Conor McHugh (2012) defends the same account of the value of true belief. He also recognises that this account can solve the triviality objection, as I will explain later. I actually think that truth is the only evaluative standard of belief, but there is no need to defend this stronger thesis here.

My account is that being good *qua* belief is epistemically valuable, because the evaluative standard of belief is *truth-directed*: a belief is good *qua* belief if *true*. There are other evaluative kinds whose evaluative standards are also truth-directed: a perception is good *qua* perception if *veridical* or *accurate enough (containing enough true information)*, an intellectual trait is good *qua* intellectual trait if the person of the intellectual trait is characteristically motivated by *love of truth* and *reliably successful in gaining true beliefs*. Similarly, a thermometer is good *qua* thermometer if its temperature measurement is *reliable* and *accurate enough*, a database is good *qua* database if its data are *largely true* and *maintained well*. Those goodness standards are truth-directed: roughly, they are good *by being true*, *by promoting truth*, *by loving truth*. Let me call an evaluative kind whose evaluative standards are truth-directed *an epistemic kind*. Epistemic values, as I define them, are *goodness of epistemic kind*. My account of epistemic value as goodness of epistemic kind can, therefore, satisfy (D1). It offers an explanation as to why the value of true belief *qua* belief is an epistemic value rather than aesthetic value or practical value, because plausibly the evaluative standards for things of aesthetic value or practical value are not truth-directed, but rather directed at, say, beauty or well-being.

One may respond that my account is *ad hoc*, namely, that my account of epistemic kinds that their goodness standards are *truth*-directed may seem arbitrary. But one should notice that it is a fact that there are evaluative kinds that their goodness standards are truth-directed. I only stipulate that those are *epistemic* kinds. This seems plausible enough, because it is hard to deny that truth is central to the epistemic.

Hence, my account of epistemic value can satisfy all of the desiderata of an account of epistemic value (I will discuss the consequential problem in the next section). Epistemologists, therefore, should prefer my account to the orthodox view. If my account is accepted, the research on epistemic value would veer into a new, more fruitful direction. Let me explain how.

As I discuss in the introduction, the orthodox view tends to lead epistemologists into debating what the epistemic goal is. For example, we have seen the triviality objection that truth is the only epistemic goal. Some (Ahlstrom-Vij and Grimm 2013; David 2001; 2014; Grimm 2008; Pritchard 2014) nevertheless defend veritism. Some (Smithies 2012; Williamson 2000) instead argue that knowledge is the epistemic goal. Some (Kvanvig 2014) argue for pluralism about epistemic goals. Since the evaluative standards of epistemic kinds are truth-directed, we could

say, in some sense, that truth is the epistemic goal. But my account is different from veritism, because my account does not regard true belief as the only thing epistemically valuable for its own sake. Instead of debating about what the epistemic goal is, my account turns epistemologists' attention to more substantial issues: what are good cognitive faculties? What are good intellectual characteristics? What are good scientific tools? What is a good learning environment? I think that those questions are epistemically more important than the original question concerning the epistemic goal.[21]

This observation is connected with the next two points. First, my account also offers epistemologists guidance about how to study epistemic values. Since epistemic values are goodness of epistemic kinds, epistemologists should examine these questions: 'What are epistemic kinds?' and 'What are their evaluative standards?' For example, textbook is an epistemic kind. But how do I know what are the evaluative standards of textbook? I consider what the function of textbook is. Presumably, the functions of textbook is to allow the readers to acquire truths about a certain subject matter, so we can determine what the evaluative standards of textbook are: such as, whether it contains enough truths about the subject matter, whether it is easy for readers to learn those truths, whether it can arouse the interests of the readers to learn those truths. So my account offers a better guidance as to assess epistemic values and settling disputes about epistemic values than the orthodox view. As we have seen, the orthodox view only leads us into an unsolvable conflict between intuitions.

Moreover, although my account is an account of *epistemic value*, my account is actually discarding this jargon. The orthodox view invents the concept of epistemic value as a technical term: epistemic values are values from the epistemic point of view. It encourages talks of moral values as values from the moral point of view, aesthetic values as values from the aesthetic point of view, prudential values as values from the prudential point of view. This implies that those values may not have any value outside a peculiar point of view. But this is not true. A true belief is valuable for being a good belief, a beautiful artwork is valuable for being a good artwork, a benevolent person is valuable for being a good person. They are not

[21] In the original debate on the epistemic goal, the epistemic goal is often understood in terms of the goal of inquiry or cognition. However, it is not true that if the goal of inquiry is, say, truth, then one must regard truth as the epistemic *intrinsic* value. For inquiry and cognition are also, arguably, epistemic kinds. Hence, even if the goal of inquiry and cognition is truth, that only means that an inquiry or a cognition that satisfies the truth goal is good *qua* inquiry or cognition. Hence, the claim that truth is the goal of inquiry or cognition does not necessarily support veritism that truth is an epistemic intrinsic value.

valuable only from a certain point of view, and we philosophers are not interested in them merely because they are valuable from a certain point of view, we are studying them from a certain point of view. Rather, we are interested in what makes beliefs, artworks, or people good, and we study them as what they are. What we are studying are *values in epistemology, ethics,and aesthetics*.

I have argued that my account has several advantages over the orthodox view. But I have not shown how my account can solve the consequentialist problem. That is my goal in the next section.

1.5 Truth-conduciveness without Teleology

In this section, I will first explain why epistemic teleology suffers from the consequentialist problem and then show how my account can solve it.

A teleological view, as Selim Berker (2013a; b) explains, contains three components: (i) *a theory of intrinsic value* that specifies what are intrinsic values; (ii) *a theory of overall value* that teaches how to measure the overall intrinsic value of certain things and to rank them accordingly; and (iii) *a deontic theory* that determines the deontic properties, such as being right, justified, or permissible, to things on the ground of the theory of overall value. The theory of intrinsic value is the most fundamental one, on which the other two theories depend. For a theory of epistemic teleology, it is paramount to have a theory of epistemic intrinsic value. Assuming that truth is the only epistemic intrinsic value, epistemic teleology will explain epistemic justification in terms of *truth-conduciveness*. As Berker (2013a; b; 2015) observes, epistemic teleology is pervasive in contemporary epistemology: 'in many circles, the predicate "… is a source of epistemic justification" has almost become synonymous with the predicate "… is truth-conducive"' (2013a, 383). A paradigm example of epistemic teleology is reliabilism. Berker claims that 'a more accurate name for reliabilism is "truth-conducivism"' (2013b, 363).

Although the account of epistemic justification in terms of truth-conduciveness is prevalent and looks innocuous, Berker (2013a; b; 2015) argues that epistemic teleology is seriously flawed, though his reason has changed over time. In Berker (2013a), his reason is that epistemic teleology, like its ethical counterpart, admits objectionable trade-offs: both accept small sacrifices (killing one person, having one false belief) in order to achieve greater good (saving the lives of five people, obtaining many true beliefs). However, in Berker (2013b, 2015),

he holds that there is a more deeper structural problem underlying trade-offs that renders epistemic teleology problematic, though Berker claims that he is still struggling to figure out what the structural problem is. I would like to suggest that the underlying structural problem is simply a familiar one about teleology, i.e., *only consequence matters*. For teleology regards only certain kinds of consequence as intrinsically valuable, actions or rules that govern actions are only instrumentally valuable for being effective in promoting the intrinsic values. Epistemic teleology must explain epistemic justification *only with regards to the consequences* it can bring about: namely, the proportion and number of true beliefs over false beliefs. As the problem of objectionable trade-offs shows, however, we seem to think that *the way* justification is conducive to truth does matter to epistemic evaluation. Consequence is not the only factor when we consider epistemic justification. But epistemic teleology cannot explain why the way of truth-conduciveness matters to epistemic evaluation. This is what I call the *consequentialist problem* for epistemic teleology.

Let me explain the consequentialist problem through one of Berker's counterexamples to epistemic teleology. In this example, Berker's target is process reliabilism:

> Suppose the following is true of me: whenever I contemplate whether a given natural number is prime, I form a belief that it is not. "Is 25 prime? No, it is not." "Is 604 prime? No, it is not." "Is 7 prime? No, it is not." Let us also stipulate that this is the only cognitive process by which I form beliefs about the primeness of natural numbers. (I'm a simpleminded kind of guy.) Since the ratio of prime to composite numbers less than n approaches 0 as n approaches infinity, my belief-forming process tends to yield a ratio of true to false beliefs that approaches 1. Therefore process reliabilists are forced to say that, because my belief-forming process is almost perfectly reliable, any belief formed on its basis is justified. But that's crazy! (Berker 2013b, 374)

Berker regards that this is a counterexample to process reliabilism. More importantly, it is the structure of process reliabilism that gives rise to this counterexample, since the reliability of the process is determined entirely by the ratio of true beliefs over false beliefs that it can produce. In Berker's prime number example, the process of generating beliefs about natural

numbers being not a prime number is highly reliable. Therefore, process reliabilism has no choice but to give the wrong verdict that beliefs thus produced are justified.[22]

I have shown that the consequentialist problem is a structural problem of epistemic teleology. Now I will explain why my account does not face the consequentialist problem. To avoid the consequentialist problem, we need not give up the idea that justification is a means to true beliefs. The idea that justification is a means to truth seems to mislead epistemologists into becoming teleologists. As Laurence Bonjour writes:

> What makes us cognitive beings at all is our capacity for belief, and the goal of our distinctively cognitive endeavours is *truth*: we want our beliefs to correctly and accurately depict the world. . . . The basic role of justification is that of a *means* to truth, a more directly attainable mediating link between our subjective starting point and our objective goal. . . . If epistemic justification were not conducive to truth in this way, if finding epistemically justified beliefs did not substantially increase the likelihood of finding true ones, then epistemic justification would be irrelevant to our main cognitive goal and of dubious worth. . . . Epistemic justification is therefore in the final analysis only an instrumental value, not an intrinsic one. (Bonjour 1985, 7-8)

As I have argued, even if truth is the epistemic goal, that does not imply that truth must be regarded as the epistemic intrinsic value. Even a means can be evaluated *qua* its own kind. The fact that a car is a means to transportation does not mean that its value is exhausted by its instrumental value to transportation. A car can be evaluated *qua* car, which could be finally valuable.[23] Hence, although epistemic teleology and my account both regard justification as a means to truth, our difference is as follows. Epistemic teleology evaluates justification entirely in terms of the ratio of true beliefs over false beliefs that it can produce. It does not matter how those beliefs are produced. Hence, the process in Berker's prime number example would be counted by epistemic teleologists as reliable.

---

[22] Kristoffer Ahlstrom-Vij and Jeffery Dunn (2014), on the contrary, bite the bullet and argue that process reliabilism is right that those beliefs are justified. Besides, Alvin Goldman (2015) argues that Berker misinterprets process reliabilism and his version process reliabilism is immune to Berker's counterexamples. However, Goldman's responses seem to me unconvincing. See Berker's reply (2015) to Goldman.

[23] This point is argued by Shelly Kagan (1998).

However, my account evaluates justification in terms of goodness of epistemic kind. The evaluative standards are not determined *entirely* by the consequence. Instead, the standards are determined by the epistemic kind itself. Different theories of justification would have different candidates for justification. To see how my account is applied to epistemic justification, let us see an example. I believe that Alvin Plantinga's proper functionalism (1993) is a non-teleological version of reliabilism, to which my theory of epistemic value is applicable.[24] Here is Plantinga's account of epistemic justification (though he uses the term 'warrant' instead):

> A belief has warrant for me only if (1) it has been produced in me by cognitive faculties that are working properly (functioning as they ought to, subject to no cognitive dysfunction) in a cognitive environment that is appropriate for my kinds of cognitive faculties, (2) the segment of the design plan governing the production of that belief is aimed at the production of true beliefs, and (3) there is a high statistical probability that a belief produced under those conditions will be true. (Plantinga 1993, 46-47)

Although the third condition in Plantinga's theory looks teleological, the first two conditions offer non-teleological elements: cognitive faculties are epistemic kinds that function to produce true beliefs. A good cognitive faculty must not only be truth-conducive (the third condition), but also be truth-conducive *in ways proper to its function or design*. For example, suppose that whenever I see a tree, my perceptual belief that it is a tree would always induce me to believe '*n* is not a prime number' (where *n* is a natural number). But my perception can only provide justification for my perceptual belief about trees, but not my belief about prime numbers. The reason is that the function of my perceptual faculty is to produce beliefs about visible features of macro-objects in suitable environmental conditions. It does not function to produce beliefs about prime numbers. Even if the production of the beliefs about prime numbers is more truth-conducive than that of the perceptual beliefs, the former still lacks justification.

[24] Also, I believe that virtue epistemology or agent reliabilism (Greco 2010; Sosa 2007; Zagzebski 1996) also offers a non-teleological account of epistemic justification. Berker himself also recognises this point, as he observes: 'I think there is a way of reading the evolution of Ernest Sosa's brand of virtue epistemology according to which he has gradually shifted from a consequentialist conception of epistemic virtues in terms of truth-conduciveness to a more truly Aristotelian way of conceiving of the epistemic virtues' (Berker 2013b, 381)

I have shown that my account offers a solution to the consequentialist problem. The consequentialist problem shows that justification cannot be simply truth-conduciveness. My account explains justification instead in terms of *truth-conduciveness in ways proper to a certain epistemic kind*. However, let me consider an objection. One may object that my account is not better than epistemic teleology, because my account cannot explain away Berker's prime number example. The argument is that one could stipulate the belief-forming process about prime numbers as an epistemic kind and its standard as the ratio of true beliefs over false beliefs that it can produce. So the argument concludes that my account would still consider those beliefs generated by the process as justified. My account is not better than epistemic teleology.

To respond to the objection, I want to emphasise that my aim here is not to provide an account of epistemic justification. I do not claim that a belief is justified if a belief is produced in a way proper to an epistemic kind (whatever it is). My point here is a modest one: my account does not bring the consequentialist problem epistemic teleology has. Epistemic teleology must regard the belief-forming process about prime numbers as able to provide justification, because epistemic teleology holds that truth is the intrinsic value and whatever can promote truth can provide justification. Instead, I suggest that justification can be explained in terms of goodness of epistemic kinds, but I do not mean that any goodness of epistemic kinds can provide justification. What sort of goodness of epistemic kind can provide justification is another task that I cannot deal with here. Nevertheless, even if the process in Berker's example can be considered an epistemic kind, I am not forced to hold that those beliefs are justified, which is not the case for epistemic teleologists.

To recap, epistemologists are right to take truth-conduciveness as the hallmark of epistemic justification. But they are wrong to regard the relation between truth and justification as the one between intrinsic value and instrumental value, which misleads them to teleology. My theory of epistemic value offers a non-teleological account of epistemic justification. Epistemologists can keep regarding justification as the means to true beliefs, but evaluate the means as its own kind that has its own standards.

1.6 Conclusion

I have shown how the orthodox view—that epistemic value is an intrinsic value from the epistemic point of view—gives rise to the triviality objection and the consequentialist problem.

I offer a novel account of epistemic value as goodness of epistemic kinds and show how my account satisfies all of the desiderata. However, I can imagine that an epistemic teleologist would respond that I miss the point of why epistemic value is considered as intrinsic value. Indeed, he may argue that (D2) should be that an adequate account of epistemic value must show that there are epistemic intrinsic values rather than epistemic final values. Since not all final values are intrinsic values, my account fails to satisfy (D2).

But why should epistemic value be a kind of intrinsic value? Proponents of the orthodox view would respond that a theory of epistemic evaluation should be able to explain why epistemic normativity is categorical: that is, why there are *pro tanto* reasons, *period*, to comply with epistemic norms, such as we should not have false or unjustified beliefs. A theory of epistemic evaluation based on my account cannot explain why epistemic normativity is categorical, because goodness of a kind does not provide categorical reasons to pursue it. For example, the fact that the coffee machine is good *qua* coffee machine does not offer me a categorical reason to have it. At most, what it can offer is a conditional reason to have the coffee machine *if* I want a good coffee machine or I should have a good coffee machine. Similarly, the fact that a true belief that p is good *qua* belief could only provide a conditional reason to have that belief if we want a good belief about p or we should have a good belief about p. However, intrinsic value, as shown in the isolation test INTENTIONAL, is able to provide categorical reason to pursue it. Hence, epistemic value should be a kind of intrinsic value.

However, although I agree that an account of epistemic value ought to provide the basis of an account of epistemic normativity, the orthodox view is wrong that appealing to epistemic intrinsic value could do that. I will give my full account in Chapter 5, but let me briefly explain why the orthodox view is wrong. The reason is simple, because an epistemic intrinsic value is intrinsically valuable *merely* from the epistemic point of view. The reason it provides is still a conditional reason to pursue it if there are further reasons to view things from the epistemic point of view. For example, since truth, according to the orthodox view, is intrinsically valuable from the epistemic point of view, the orthodox view does not explain why we should not have false beliefs, *period*. For we may think that some false beliefs are intrinsically valuable, *period*, even though we agree that all true beliefs are intrinsically valuable from the epistemic point of view. Hence, the orthodox view is not better than my account.

It is right that intrinsic value can, but goodness of a kind cannot, provide *pro tanto* reasons to pursue it. But that does not mean that some goodness of a kind is not intrinsically valuable.

Indeed, in Chapter 5, I will argue that some goodness of epistemic kinds is intrinsically valuable, so that my account is able to explain epistemic normativity. However, my account of epistemic normativity is also built on a virtue-theoretic account of epistemic evaluation, which I will develop in the next three chapters.

Chapter 2.

Epistemic Evaluation as Performance Evaluation

2.1 Introduction

Virtue epistemology has emerged as a central force in contemporary epistemology. Virtue epistemologists often claim that the virtue-theoretic account enjoys several advantages. First, virtue epistemology can solve the Gettier problem (Zagzebski 1996, 283-99; Sosa 2007; Greco 2010, ch.5). Second, virtue epistemology can solve the so-called value problems of knowledge (Riggs 2002; Zagzebski 2003; Sosa 2007, ch.4; Greco 2010, ch.6). Third, virtue epistemologists (Greco 2010; Sosa 2007; 2011) offers a more unified, elegant account of normativity, because they argue that epistemic normativity is merely a species of performance normativity (which I will explain later).[25]

In this chapter, I will discuss the last supposed advantage the virtue-theoretic account claims to enjoy. The third advantage is crucial, because whether the virtue-theoretic account can enjoy the first two advantages largely depends on the third one. I will postpone the discussion of the virtue-theoretic account of the value problems of knowledge to the next two chapters. (In the appendix of Chapter 4, I also briefly discuss how the virtue-theoretic account could respond to the Gettier problem.) Before that, we need to properly understand the claim that epistemic normativity is a species of performance normativity.

Here are some clear statements by John Greco and Ernest Sosa that epistemic normativity is a species of performance normativity:

[25] Linda Zagzebski (1996), instead, regards epistemic normativity as a species of ethical normativity and beliefs as a kind of action. Despite the difference, the structure of Zagzebski's account is similar to the account of epistemic normativity as a species of performance normativity by Ernest Sosa and John Greco. However, I think Zagzebski's view is less plausible (see Heather Battaly (2008) for a comparison of Zagzebski's view and the view of Ernest Sosa and John Greco). So I choose to focus on Sosa and Greco's account here.

Epistemic Evaluation as Performance Evaluation

> Knowledge is a kind of *achievement*, or a kind of success for which the knower deserves credit. And in general, success from ability (i.e. achievement) has special value and deserves a special sort of credit. This is a ubiquitous and perfectly familiar sort of normativity. Thus we credit people for their athletic achievements, for their artistic achievements, and for their moral achievements. We also credit people for their intellectual achievements. Epistemic normativity is an instance of a more general, familiar kind. (Greco 2010, 7)

> Belief is a kind of performance, which attains one level of success if it is true (or accurate), a second level if it is competent (or adroit), and a third if its truth manifests the believer's competence (i.e., if it is apt). Knowledge on one level (the animal level) is apt belief. The epistemic normativity constitutive of such knowledge is thus a kind of performance normativity. (Sosa 2011, 1)

Although the terminologies are different, their account of knowledge is basically the same. They both hold that knowledge is a belief's being true because of the exercise of the believer's abilities (which should be competent enough) and that the evaluation of belief shares the same structure of the evaluation of performance. In this chapter, I will focus particularly on Sosa's view, but my discussion is also applicable to Greco's view.

Sosa holds that belief is a kind of performance, because *beliefs aim at truth*, and *all performances with an aim* can be evaluated under performance evaluation:

> Performances with an aim, in any case, admit assessment in respect of our three attainments: accuracy: reaching the aim; adroitness: manifesting skill or competence; and aptness: reaching the aim *through* the adroitness manifest. (Sosa 2007, 23)

> If beliefs constitute (such) performances, they must each have an aim. ... Belief is always evaluable for truth (positively) or falsehood (negatively), after all, and any performance positively evaluable by whether it attains a certain status may be said, analytically, to 'aim' for that status. (Sosa 2011, 15)

Sosa holds, then, that knowledge is apt belief, that is, 'an apt epistemic performance, one that manifests the relevant competence of the believer in attaining the truth' (Sosa 2011, 4). Hence, epistemic evaluation is a species of performance evaluation. In other words, the virtue-theoretic account holds that the way we evaluate belief is essentially the same as the way we

evaluate our performance. A performance can be evaluated in terms of whether it successfully fulfils its aim (accurate, successful), whether it is skilfully or virtuously executed (adroit, competent, skilful), and whether it successfully fulfils its aim because it is skilfully or virtuously executed (apt). Knowledge is a successful belief because of being virtuously-produced.

To that extent, Sosa's virtue-theoretic approach appears to overthrow the orthodox view, at least in the analytic tradition, that belief and knowledge are mental states. Is the virtue-theoretic approach correct? One should notice that the virtue-theoretic account has not offered any argument that belief is a kind of performance rather than a kind of mental state. The virtue-theoretic account is supported by the observation that epistemic evaluation is structurally similar to performance evaluation and the argument that the theoretical advantages that modelling epistemic normativity on performance normativity can solve several problems in epistemology. As Matthew Chrisman (2012) objects, however, since the virtue-theoretic account has not proved that belief is a kind of performance, it may simply change the subject matter. Indeed, Chrisman argues that belief and knowledge are states rather than performances, so that epistemic evaluation cannot be modelled on performance evaluation.[26] Despite the theoretical advantages of virtue epistemology, Chrisman concludes that it should be rejected.

The aim of this paper is to critically engage with the debate between Chrisman and the virtue-theoretic account. In section 2, I discuss Chrisman's argument for the thesis that, since belief and knowledge are states, performance evaluation is not applicable to belief. The virtue-theoretic account may respond that epistemic evaluation can be modelled on performance evaluation, because maintaining (and forming) a belief is a performance. Chrisman responds that maintaining a belief is not a performance, but I argue that his response is wrong. However, the thesis that maintaining a belief is a performance cannot really save the virtue-theoretic account. The reason is that belief and knowledge are states rather than performances. The fact that performance evaluation is applicable to maintaining a belief does not show how a state, such as belief or knowledge, can be evaluated as a performance. In section 3, I discuss Sosa's view, as seen in the above quotations, that a state with an aim can be regarded as a performance for being assessable under the structure of performance evaluation. However, I argue that Sosa's view should be modified, because although it is true

[26] The same objection can be found in Pascal Engel (2013a; b) and Kieran Setiya (2013). I focus on Chrisman's argument because his argument is the most developed one.

that belief can be evaluated under performance evaluation, when a belief is assessed as competent or apt, what is competent or apt, *strictly speaking*, is not belief itself but the cognitive performance that produces the true belief. So what is *primarily* evaluated under performance evaluation is the cognitive performance, and belief is evaluated under the structure of performance evaluation *only derivatively*. If whatever falls under performance evaluation is a performance, then belief is a performance, but only in the derivative sense.

The thesis that belief is a performance only in the derivative sense is crucial to the virtue-theoretic account. First, the virtue-theoretic account can maintain that performance evaluation is applicable to things that are not performances, such as states (e.g. belief) or objects (e.g. paintings). Second, the thesis has important implications for the virtue-theoretic account of the value of knowledge. Knowledge, according to the virtue-theoretic account, is an apt (true because adroit) belief. Since knowledge, as a state, is a performance only in the derivative sense, what is adroit and apt is not the true belief itself, but the cognitive performance. Therefore, there are two sources of the value of knowledge: *the value of true belief* and *the value of adroit cognitive performance*. Most epistemologists try to explain the value of knowledge solely in terms of valuable properties of belief, because they think that knowledge is a kind of belief. In the subsequent chapters, I will argue that this reasoning is the root of the problems about the value of knowledge. My discussion of the virtue-theoretic account shows that this reasoning is flawed. The fact that knowledge is a kind of belief does not mean that its value must be explained solely in terms of the values of properties of belief. Hence, this chapter paves the way to my discussion of the value problems in the subsequent chapters.

2.2 Chrisman's Linguistic Objection: 'Believe' and 'Know' are Statives

In this section, I will present Chrisman's objection (2012) to the virtue-theoretic approach that belief is a state, not a performance. As Chrisman argues, there is solid linguistic evidence showing that belief and knowledge should be classified as states rather than performances. English verb phrases can be categorised in terms of *aspect*: *statives* and *non-statives*. Statives are verb phrases referring to *states*, such as 'He wants the pie'. Non-statives are verb phrases referring to *occurrences* or *events*, such as 'She is painting the portrait'. Chrisman argues that verb phrases containing 'believe and 'know' should be categorised as statives, so belief and knowledge are states rather than performances.

Some linguist tests can reliably distinguish whether a verb is a stative or non-stative. First, non-statives can have the progressive form of verbs, while statives cannot. Using Sosa's favourite analogy between archery and belief (Sosa 2003; 2007; 2011), which attempts to indicate that belief is a kind of performance:

(1a) Archie was shooting the arrow skilfully;

(1b) The arrow is hitting the bull's-eye;

(2a) Stephen Hawking was believing that black holes exist;

(2b) Stephen Hawking is knowing that black holes exist.

(1a) and (1b) are perfectly legitimate English sentences, but (2a) and (2b) are not. A relevant test is that when non-statives are used in the simple present form, such as 'Archie shoots the arrow', it implies that the occurrences or events are repetitive, serial, or habitual, while statives so used, like (2a) and (2b), do not have the same implication. These linguistic tests show that 'shoot' and 'hit' are non-statives, whereas 'believe' and 'know' are statives.[27] Therefore, Chrisman concludes that Sosa cannot use the archery analogy to show that belief and knowledge are performance. Belief and knowledge, hence, cannot be evaluated in the same way as performance. Epistemic evaluation is not a species of performance evaluation.

The linguistic evidence provided by Chrisman, to my mind, conclusively demonstrates that belief and knowledge are states of a cognitive subject. Chrisman's argument shows that Sosa's archery analogy cannot support the view that belief is a performance, on the grounds that 'shoot' is a non-stative and 'believe' is a stative. But it has not proved that states cannot be performances. Indeed, Sosa does not deny that belief is a state, but he maintains that some states can be performances:

> Some acts are performance, of course, but so are some sustained states. Think of those live motionless statues that one sees at tourist sites. ... Beliefs too might thus count as performances, long-sustained ones, with no more conscious or intentional an aim than that of a heartbeat. (Sosa 2007, 23)

[27] Chrisman also mentions that occasionally some statives can be used in the progressive form, such as 'We are living in Spain' and 'Southampton FC are lying second in the league'. Chrisman explains that linguists usually explain this anomaly as expressing that the states are temporary, but admits that a deeper explanation may be required. However, he thinks that the linguistic data are enough to show that belief and knowledge are states.

Epistemic Evaluation as Performance Evaluation

Sosa uses the analogy of a live motionless statue to support that a long-maintained state, of which belief is an example, is a kind of performance. Unfortunately, this analogy does not work as well, because while *being* a live motionless statue is a state, it is not a performance. The performance is *maintaining a live motionless statue*. So the analogy cannot show that performance evaluation is applicable to states.

Perhaps, what Sosa means by the statue analogy is that epistemic evaluation can be modelled on performance evaluation, because what is under epistemic evaluation is performances like maintaining beliefs.[28] It seems clear that maintaining a live motionless statue is a performance, since 'maintain' is a non-stative (because using 'maintain' in the progressive form is legitimate in English.) Accordingly, if belief is not a performance because 'believe' is a stative, Sosa could hold that maintaining a belief is a performance.

Chrisman agrees that 'maintain' is a non-stative. But he argues that maintaining a live motionless statue is not strictly a kind of performance, so performance normativity cannot be applied to it. Chrisman's argument from (2012, section IV) is that non-statives can be distinguished into three kinds: *process*, *development*, and *punctual occurrence*. The distinction is based on two features of verbs: *duration* and *telicity* (that is, a telic verb phrase implies that the action or event takes time *with* a specific endpoint, and an atelic verb phrase implies that the action or event takes time *without* a specific endpoint). One can distinguish these kinds of non-stative by combining the past progressive and perfect tense of a verb in a single sentence. For example,

(3a) Process (durative and atelic): He was running, so he has run;

(3b) Development (durative and telic): He was filming a movie, so he has filmed a movie;

(3c) Punctual Occurrence (non-durative and telic): He was winning the race, so he has won the race.

[28] One may also include 'forming beliefs' (or others, such as 'revising beliefs', 'suspending beliefs'). Chrisman also discusses this possibility, but he argues (2012, 602-604) that this move does not save the virtue-theoretic account. I agree, but for different reasons. I think that the virtue-theoretic account eventually needs to show how performance evaluation is applicable to states, since knowledge is a state. To avoid unnecessary complication, I will not discuss the formation of beliefs here.

(3a) sounds natural in English, since if someone has run, it is correct to say that she was running. This shows that the verb, 'run', is durative and atelic. In contrast, (3b) sounds odd. The reason is that someone who was filming a movie may never finish it, which shows that the verb, 'film', is durative and telic. Whether (3c) is right is unclear. If (3c) is a legitimate English sentence, the verb, 'win', is telic. But when the phrase, 'is winning the race', is used to imply an endpoint, this phrase refers to the point of reaching the endpoint, rather than to the process of reaching it. Therefore, 'win' is non-durative.

The point of differentiating these kinds of non-stative, Chrisman argues, is that paradigmatic verbs of performance used by Sosa, such as 'shoot' and 'hit', are classified as development or punctual occurrence, while 'maintain' is classified as process:

(4a) Process: He was maintaining the belief, so he has maintained the belief;

(4b) Development: Archie was shooting the arrow, so he has shot the arrow;

(4c) Punctual Occurrence: The arrow was hitting the target, so it has hit the target.

(4a) is as natural as (3a), which shows that 'maintain', like 'run', is durative and *atelic*, so that it should be classified as process. However, (4b) is more like (3b), and (4c) more like (3c). This shows that the verbs 'shoot' and 'hit' are classified as development or punctual occurrence (they are both *telic*). Given their difference in telicity, Chrisman argues that virtue epistemologists cannot hold that maintaining a live motionless statue or maintaining a belief is a performance, because paradigmatic performances are expressed by telic verbs.

However, Chrisman's argument is too hasty. It relies on the claim that verb phrases of paradigmatic performances are telic rather than atelic. However, that claim is simply wrong. Many verb phrases that express some kinds of performance are atelic. Singing is certainly a paradigmatic kind of performance, and there is nothing wrong to say, 'She was singing, so she has sung'. Chrisman may respond that singing by itself is not a performance. Instead, singing *Yesterday* is a performance, because it may be interrupted so that one does not finish singing the song. Chrisman could thus argue that 'is singing *Yesterday*' is telic (a development non-stative). Accordingly, Chrisman would hold that it is odd to say 'He was singing *Yesterday*, so he has sung *Yesterday*'. However, it seems wrong that one has to finish singing a song, in order to be counted as having sung that song. That sentence sounds natural to me, even if I know that the singer did not finish that song. So that suggests that the phrase 'is singing *Yesterday*' is atelic, and yet singing *Yesterday* is a performance.

More importantly, even granted that 'is singing *Yesterday'* is a development non-stative, it does not show that 'is singing' is not a performance. One can certainly improvise without knowing how it will end. It is hard to believe that improvisational singing is not a kind of performance.

Chrisman's mistake, I think, is that he argues from the fact that the verb phrases of *paradigmatic* performances are telic to the conclusion that *only* those expressed by telic verb phrases are performances. But singing is no doubt a performance, as is maintaining a live motionless statue. People certainly regard those live statues as performers and are willing to pay them to have photos with them. Surely, singing and maintaining a live motionless statue are *paradigmatic* performances. Hence, Chrisman is simply making an illicit generalisation.

To be fair, Chrisman does try to explain why maintaining a live motionless statue is not a performance. However, his explanation is not plausible:

> What about those live motionless statues? Their performance, in my view, is maintaining a pose for some period of time. It requires no great skill to stand still for 10 seconds, but doing so for 10 minutes might be a performance worth a tip. However, one can maintain a belief for 10 seconds or 10 minutes or for any length of time, and the belief can count as knowledge, no matter how long one maintains it. So, if our ordinary ways of talking about these things are indicative or our ordinary conception of these things, which I think they are, then, on our ordinary conception of maintaining a belief, this is a process (Vendler's term: activity) rather than an event (Kenny's term: performance). But, *since we use event-descriptions rather than process-descriptions to talk about all paradigmatic performances, maintaining a belief is not plausibly thought of as a performance.* Again this means that the structures of performance normativity do not apply to it. This is not just an accident of grammar. The *telic aspect of performances* seems to be part of why we can distinguish along the three dimensions of performance normativity: success, skill, and aptness. Hence, something *atelic*, such as maintaining a belief, will never be the right sort of thing to which to apply these evaluative distinctions. (Chrisman 2012, section IV; my italics)

Note that Chrisman's sole reason why performance normativity cannot be applied to maintaining belief, in this quotation and the argument discussed above, is that maintaining belief is *atelic*. But maintaining a live motionless statue appears to be a performance. Chrisman

tries to explain that maintaining a pose itself is not a performance, but only maintaining a pose for a long time is a performance, because the latter needs *great skill*. But Chrisman does not explain why skilfulness determines what counts as a performance. It is true that maintaining a live motionless statue for a long time is a good performance of maintaining a live statue, but that does not mean maintaining the statue for a few minutes is not a performance at all. Maintaining the statue for a few minutes is a bad performance, but still a performance. Chrisman's argument would imply that any unskilful atelic performance is not really a performance. But that is just wrong. I can certainly sing *Yesterday*, though my singing would be terribly off-key and I can't even sing the lyrics correctly. My singing is horrible, but still a performance. Hence, Chrisman's argument that maintaining a live motionless statue is not a performance is unsuccessful.

Furthermore, Chrisman holds that the structure of performance evaluation is essentially shaped by the telicity of performance, so performance evaluation is inapplicable to atelic actions or events. However, I cannot see how the atelicity of maintaining a pose makes performance evaluation inapplicable. Undoubtedly, a telic performance can be counted as successful when it reaches the endpoint implied by its telicity. Hence filming a movie or hitting the target can be evaluated as successful if the movie has been produced or the target has been hit. But it does not follow that an atelic action or event cannot be evaluated in the structure of performance evaluation. True, an atelic action does not have a specified endpoint, so that its success does not depend on reaching the endpoint. On the contrary, the fact that a process non-stative is durative and atelic (lasting without a specified endpoint) means that the success of an atelic performance could be determined by the constant continuation of the process. Hence, in the case of maintaining a motionless pose, one can evaluate whether the performance is successful (how long the performer maintains a motionless pose), whether the performance is competent at maintaining the motionless pose, and whether the performance is apt (whether the success of performance is attributable to the competence of the performer). Or consider improvisational singing. Improvisational singing is also atelic, but its success may not be determined merely by the continuation of singing. A brief improvisational singing could also be successful, if it is tuneful, original, or entertaining. And no doubt improvisational singing can be assessed in terms of adroitness and aptness.

So we see that Chrisman's linguistic objection fails to show that performance evaluation cannot be applied to belief maintenance, because he is wrong to hold that atelic actions or events are not performances. However, that does not save the virtue-theoretic account from

Chrisman's objection. What the virtue-theoretic account holds is that belief, which is a state, can be evaluated under the structure of performance evaluation. Chrisman has convincingly shown that belief is a state. So the fact that performance evaluation is applicable to maintaining beliefs does not mean that it is applicable to belief. In the next section, I will discuss whether performance evaluation is applicable to states.

2.3 Chrisman's Objection (II): Is a State a Performance?

We have seen that Chrisman's objection to the virtue-theoretic account is that belief and knowledge are significantly different from paradigmatic performances, such as archery, in that shooting is an event and believing is a state. But Sosa's archery analogy is not meant to show that believing is literally an action or an event (as we have seen, Sosa still regards beliefs as states). In Sosa's recent article, he distinguishes two kinds of performance: '*functionings* (functionally assessable states) and '*endeavours* (with a freely determined aim)' (2013, 585). More roughly, the reason why functioning and endeavour are performances is that both have certain *aims*. As Sosa explains:

> One might alternatively opt for a broader notion of 'endeavour' according to which any pursuit of an aim, any teleology even if merely functional, would involve 'endeavoring' to attain that aim, as does the heart when it beats regularly, aiming to circulate the blood. (2013, 585)

> Since belief is a functional state and aims at truth, belief is a kind of performance, according to Sosa' view. Therefore, Sosa would respond to Chrisman that, when determining whether something is a performance, 'the distinction between the passive [statives] and the active [non-statives] is here negligible by comparison with the distinction between states that are and those that are not functionally assessable' (2013, 588).

Chrisman, however, foresees this sort of reply. Chrisman acknowledges that the virtue-theoretic approach may use the notion of performance broadly. However, he argues that this move cannot demonstrate the applicability of performance evaluation to belief:

> One might grant that neither beliefs nor maintaining beliefs are performances but argue that Sosa's theoretical structure of 'performance normativity' applies more generally than just to performances. … Perhaps anything with an aim could be

> said to admit of this distinction by analogical extension. Then, even if beliefs are states rather than performances, as the aspectual data seem to indicate, it may still be true that beliefs, because they 'aim at the truth', admit of the normative distinctions characteristic of performances. And when a belief successfully reaches its 'aim' because of the skill of the believer, it is knowledge. The problem I have with this response is that it strikes me as completely *ad hoc*. Unless there are other states that plausibly admit of the distinctions characteristic of performance normativity, it will be entirely theory driven to think that the states of belief and knowledge fit into this normative structure. (Chrisman 2012, 605)

Chrisman argues that Sosa needs to offer more examples of state to which performance evaluation is applicable. Otherwise, what Sosa does is simply assume that performance evaluation is applicable to belief. In the previous section, we have seen that Sosa uses the example of a live motionless statue to support his view that a long-sustained state, like belief, can also be regarded as a performance. However, as we have discussed, this example only shows that performance evaluation is applicable to maintaining a live statue, which is not a state. What we need from Sosa here is an argument that performance evaluation is also applicable to states like being a live motionless statue or belief. However, we have not seen one.

Let us consider whether performance evaluation is applicable to the state of being a live motionless statue. Sosa says that a state with an aim is assessable in terms of performance evaluation: i.e., success, competence, and aptness. Suppose that the state of being a live motionless statue of the Statue of Liberty *aims* at looking like the Statue of Liberty: hence, the state of being the live Statue of Liberty is *successful* if it looks like the Statue of Liberty. But it seems not to be assessable in terms of *competence*. What can be assessed as competence (and apt) is *maintaining the state of being the live Statue of Liberty*. The reason for this is simple. A state with an aim can be evaluated in terms of success: the state is successful when it satisfies its aim. However, when we evaluate competence, what we evaluate is skilfulness, adroitness, or ability. But states cannot be skilful, adroit, or able. What is skilful or competent is *forming* or *maintaining* a state. Forming or maintaining a state requires skills or abilities and its aim is to be in the state. So forming or maintaining a state can be evaluated in terms of performance evaluation: successful if being in the state, competent if skilful or competent, and apt if being in the state because competent. Being in the state is, therefore, merely the end result of forming or maintaining a state and can be evaluated only in terms of success. Since

performance can be evaluated in terms of performance evaluation, being in a state, even by Sosa's own standard, is not a performance.

Accordingly, Sosa should not regard a belief as a performance, since a belief is also a state and is assessable only in terms of success. What really falls under performance evaluation is a cognitive performance that aims at true beliefs: hence, a cognitive performance is successful if the belief it produces is true, competent if it is skilful or adroit, and apt if the belief it produces is true because competent.

I have argued that a state is not a performance itself, but only the end result of a performance. As a result, a state is not assessable in terms of competence and aptness. To be sure, this point applies to other types of performance that do not aim at producing states. A painting is also an end result of a performance of painting. Similarly, a painting itself can be evaluated only in terms of success (depending on what aim a painting has: similarity, beauty, originality, the painter's intention), but not competence and aptness. Archery is the same. A shot is the end result of shooting, so it is assessable only in terms of success (depending on whether it hits the target). So a painting and a shot are not performances.

However, Sosa may respond that a painting and a shot are still assessable in terms of performance evaluation. For in ordinary speech, people do say that it is an adroit shot or it is a skilful painting. So it seems that an end result of a performance could still be evaluated under performance evaluation. That is no doubt true. I do not deny that when people appreciate Raphael's *The School of Athens*, they might say, 'How skilful the painting is!' But what is skilful, strictly speaking, is Raphael's painting skills, rather than the painting itself. Although the process of painting *The School of Athens* was completed several hundred years ago, we can still appreciate Raphael's skills by studying the painting itself. Sosa himself also recognises this point; he says that a shot is adroit when 'it manifests skill on the part of the archer' (2007, 22). Therefore, when we evaluate whether an end result of a performance is adroit, what we really evaluate is the competence of the performance itself. I am not claiming that one should not say that a shot is adroit, but only that what is really adroit is the shooting. One may say that the shot is adroit and apt, but only *derivatively*. If a performance can be evaluated in terms of performance evaluation, a shot can be regarded as a performance, but only in a derivative sense. What is primarily evaluated under performance evaluation is the shooting itself. So only the shooting is properly a kind of performance.

Hence, the virtue-theoretic account is compatible with the view that belief is a kind of mental state, but it can still maintain that they are performance in the derivative sense and assessable under performance evaluation. Hence, the virtue-theoretic account can still maintain its theses:

SUCCESS*: a belief is successful if and only if it is true;

COMPETENCE*: a belief is competent if and only if it is skilful or adroit;

APT*: a belief is apt if and only if it is true because competent.

I put asterisks, because belief is evaluated as a performance only in the derivative sense. When evaluating a belief under performance evaluation, what is really under evaluation is the cognitive performance:

SUCCESS: a cognitive performance is successful if and only if the belief it forms or maintains is true;

ADROIT: a cognitive performance is adroit if and only if it is skilful or competent;

APT: a cognitive performance is apt if and only if the belief it forms or maintains is true because of its being competent.

In this section, I have argued that Sosa is wrong to hold that whatever has an aim can be assessable in terms of performance evaluation. Belief aims at truth, so when evaluating a belief *qua* belief, a belief is successful *qua* belief if and only if true. But a belief *qua* belief is not competent or apt. Strictly speaking, therefore, belief is not a performance. However, that does not mean that we cannot evaluate a belief *qua* performance. The reason is that we can evaluate a performance through evaluating the end result of the performance. Belief is the end result of cognitive performance. So we can evaluate a belief *qua* performance. Hence, a belief can be considered a performance in the derivative sense. The critical point is that when evaluating a belief *qua* performance, what we really evaluate *qua* performance is the cognitive performance itself. So Chrisman is wrong to conclude that the virtue-theoretic account should be rejected, because belief is a state.

2.4 Conclusion: A Preliminary Remark on the Value of Knowledge:

I have argued that the virtue-theoretic account of epistemic evaluation could still maintain that belief is a state *and* assessable under performance evaluation. But one must bear in mind that what is really evaluated under performance evaluation is the cognitive performance rather than the belief. According to the virtue-theoretic account, knowledge is *an apt performance*. But that does not mean that the virtue-theoretic account must deny that knowledge is a kind of belief (or mental state).[29] The virtue-theoretic account can still maintain that knowledge is a mental state and it is an apt performance only in the derivative sense.

That is to say, knowledge, according to my account, is *a kind of cognitive state (true belief) which one can only be in through competence cognitive performance(s)*. Call this kind of state *achievement*.[30] An achievement is a state one can only be in through competently exercising the right kind(s) of ability. Such states are plenty. Being a chess master is an achievement that one can be in only through one's chess skills. A chess master is one who can defeat most capable chess players through one's own chess skills. A player who has not defeated an overwhelming majority of capable chess players cannot be granted the title of chess master. Thus, being a chess master is an achievement. This point is especially clear when we try to evaluate and compare two chess players, A and B. We may rank A higher than B, because A won more games than B. But we may rank B higher than A, because B defeats more chess players than A, which may suggest that the skill of B is greater than A. Furthermore, a chess master may lose the title if he loses too many games or his chess skills significantly diminish. Being a chess master is thus an achievement. The example of chess master demonstrates that

[29] The virtue-theoretic account may understand knowledge as performance in the primary sense, so that knowledge is not literally a kind of belief, but rather belief is a result of knowing. This is certainly a possibility, but I will assume, given Chrisman's linguistic evidence, that knowledge is a kind of belief (or a mental state) in the following discussion. But my point about the value of knowledge will only be strengthened if one holds that knowledge is an apt performance in the primary sense.

[30] The verb *achieve* is a stative, because its progressive form is not generally used. One may say: 'He is achieving the goal'. As mentioned in note.3, it is possible to use a stative in the progressive form. When so used, however, it signifies that the state is temporary. However, what I say about achievement as a kind of state does not hinge on this linguistic fact. For *achievement* is used as a term of art in this chapter. The idea that knowledge is a kind of achievement is from Greco (2010). I criticise Greco's idea that knowledge is a kind of achievement in Chapter 4. Nothing in my criticism there undermines my claim here, because Greco holds that achievement is an apt performance. But my view allows *inapt* ones to be achievements, so long as they are successful state one can be in *only through one's competent performance*, which does not requires that the states be successful *because* competent.

a state can still be evaluated *qua* performance, but when it is assessed as competent, what is really competent is the master's chess skill rather than the state of being the chess master.

I have argued that belief and knowledge are states, but still assessable under performance evaluation. But when they are competent and apt, it is cognitive performance, strictly speaking, that is adroit and apt. This offers a straightforward solution to the age-old problem about the value of knowledge raised by Socrates in the *Meno*. Socrates there asks the question why knowledge is more valuable than mere true belief. He asks us to consider two guides who can lead us to the city of Larissa. One guide knows the road to Larissa, but the other only has the true belief about the road. Both guides can successfully lead us to Larissa. The question then is why knowledge is more valuable than mere true belief. This is the so-called *Meno problem*.

I will tackle the Meno problem in the next chapter, but let me briefly explain my proposal. My solution to the Meno problem is simple. Since knowledge, according to my account, is a true belief one has through one's cognitive abilities, the value of knowledge consists of two distinct kinds of value: *the value of epistemic success (the value of true belief)* and *the value of epistemic competence*. Hence, it is straightforward to see why knowledge is more valuable than mere true belief.

As I will argue in next chapter, this simple solution would not be available if epistemologists do not take up the idea that the object of epistemic evaluation is the cognitive performance rather than belief itself. Epistemologists often assume that the object of epistemic evaluation is belief because knowledge is a kind of belief. But this assumption is wrong. My virtue-theoretic account does not to deny that knowledge is a kind of belief. Knowledge, according to my account, is a state of achievement. When we evaluate whether a state is an achievement, what we evaluate is the performance responsible for the state. My virtue-theoretic account helps us discover that the value of knowledge lies in the value of cognitive performance. Hence, although the result of this chapter may appear merely as a modification of the virtue-theoretic account, it is an important step towards the answers to the problems about the value of knowledge.

In conclusion, I have critically discussed Chrisman's objection to the virtue-theoretic account. I conclude that Chrisman is right that belief and knowledge are not performance in the strict sense. Contrary to the virtue-theoretic account, belief and knowledge are performances only in the derivative sense. Although it means that one can still evaluate belief *qua* performance,

what is really under performance evaluation is the cognitive performance rather than the belief. It implies that the source of the value of knowledge lies in cognitive performance. This is the lesson we need to bear in mind—as I will argue in the next chapter—if we want to offer a satisfactory account of the value of knowledge.

Chapter 3. How Not to Swamp the Value of Knowledge

3.1 Introduction

Ever since Plato, philosophers have been curious about the value of knowledge. More precisely, they want to know why knowledge is more valuable than mere true belief. In *Meno*, Socrates asked, 'Why is knowledge more valuable than mere true belief?' If we were looking for a guide who could lead us to Larissa, should we prefer a guide who knew the road to Larissa rather than a guide who did not know, but had a mere true belief about the road to Larissa? Both guides could lead us to Larissa, so it seems that it should not matter which guide we choose. This points to the conclusion that mere true belief is as good as knowledge. However, this conclusion seems counter-intuitive; knowledge appears to be more valuable than mere true belief. In the literature, it is called *the Meno problem*: *why is knowledge more valuable than mere true belief*?[31]

According to orthodoxy, knowledge is regarded as a kind of true belief. To answer the Meno problem, epistemologists naturally look for some type of property of belief that adds value to true belief, which makes knowledge more valuable *qua* belief than mere true belief. Call that type of property *the Meno property of belief*. However, the idea that there is a Meno property is at odds with another popular thesis—call it *T-monism*—that truth is the sole final epistemic value. If T-monism is true, all epistemic values, except truth, seem to have only *instrumental value*, that is, they are valuable only for the sake of producing or preserving true beliefs. If a

[31] The Meno problem, so construed, presupposes that knowledge *is* more valuable than mere true belief. Although the presupposition has intuitive support, some philosophers (e.g. Sartwell 1992) reject it. However, I will work under this presupposition and my aim is to answer the Meno problem.
Also, I shall emphasise the point that I am only interested in the *epistemic value* of knowledge. Some philosophers may answer the Meno problem by arguing that knowledge is *practically* more valuable than mere true belief (Craig 1990, 16-18; Williamson 2000, 101; Hyman 2010). Whether their solutions are correct is not my concern. My concern is how knowledge is epistemically more valuable than mere true belief.

Meno property only adds instrumental value to true belief, knowledge cannot be more valuable *qua* belief than mere true belief. For the value of truth swamps any value which is only instrumental relative to the value of truth. Given T-monism, therefore, it is difficult to see how the Meno problem is answerable. In the recent literature, this difficulty in answering the Meno problem caused by T-monism is called the *swamping problem* (Kvanvig 2003; Pritchard 2011).

In this chapter, I will argue that it is wrong to assume that to answer the swamping problem, one must find the Meno property that makes a true belief more valuable *qua* belief. Indeed, I think that this assumption renders the swamping problem intractable. People accept this assumption, because they think that knowledge is a kind of belief, so the value of knowledge must consist entirely of the values of properties of belief. Call it *the belief-based conception of epistemic evaluation*. Under the belief-based conception, the Meno problem asks: 'Why is knowledge *a more valuable kind of belief* than mere true belief?' So an account of the Meno property is essential to answering the swamping problem. I argue that the belief-based conception should be replaced with *the performance-based conception of epistemic evaluation*.[32] The performance-based conception of epistemic evaluation can accept the view that knowledge is a kind of belief. What it rejects is the assumption—embodied in the belief-based conception—that the value of knowledge consists solely of the values of properties of belief. Instead, the performance-based conception maintains that the value of knowledge consists of two distinct final values: the value of true belief and the value of good cognitive performance. Once my account is fully appreciated, it will be easy to see why the search for the Meno property is wrong and how the swamping problem can be readily solved.

In Section 2, I will explain in more detail how the swamping problem poses a problem for any attempt to solve the Meno problem. In Section 3, I show how the belief-based conception causes the swamping problem. This is shown through discussing simple reliabilism and the virtue-theoretic account with regard to the swamping problem. The virtue-theoretic account is widely considered to be the theory most likely to solve the swamping problem (Kvanvig 2003, ch.4; Fricker 2009; Pritchard 2010, ch.2; Baehr 2012). Simple reliabilism, however, is often criticised for suffering from the swamping problem (Kvanvig 2003; Riggs 2002; 2007; Zagzebski 2003). I explain how the intuition that simple reliabilism suffers from the swamping problem is

[32] I defend the performance-based conception of epistemic evaluation more comprehensively in Chapter 2.

caused by the belief-based conception. However, I argue that if one shares the intuition, one should regard the virtue-theoretic account suffers the same fate as simple reliabilism.[33] In Section 4, I argue that the virtue-theoretic account is better interpreted as adopting the performance-based conception, which is why it can solve the swamping problem. However, once the performance-based conception is adopted, even simple reliabilism can solve the swamping problem. So, it is not the case that the virtue-theoretic account enjoys any advantage over simple reliabilism with regard to the swamping problem. In Section 5, I explain the performance-based conception in greater detail. In Section 6, I will compare some recent solutions to the swamping problem with my own and show how my solution is preferable to theirs.

3.2 The Swamping Problem

The swamping problem, as Duncan Pritchard (2010, 15; 2011, 248-49) presents it, can be explained in terms of three mutually inconsistent theses:

> (S1) 'The epistemic value conferred on a belief by that belief having an epistemic property is instrumental epistemic value relative to the further epistemic good of true belief';
>
> (S2) 'If the value of X is only instrumental value relative to a further good and that good is already present, then it can confer no additional value';
>
> (S3) 'Knowledge that p is sometimes more epistemically valuable than mere true belief that p'.

(S1) follows from T-monism – the view that true belief is the only final epistemic value, so other epistemic values are instrumental relative to the value of true belief. (S3) captures our intuition that knowledge is more valuable than mere true belief. Although I think that knowledge is *always* more valuable than mere true belief, notice that (S3), as Pritchard notes,

[33] One can certainly reject the intuition or try to argue that the swamping problem is solvable under the belief-based conception, as Alvin Goldman and Erik Olsson (2009) attempt to argue that a reliable process can add value to true belief. However, my point is to show how the belief-based conception is the underlying cause of the intuition that the swamping problem is serious to reliabilism. Moreover, compared to the solution of Goldmand and Olsson, it is obvious that the performance-based conception offers a more simple and straightforward solution to the swamping problem. Indeed, once one see the advantage of the performance-based conception, one may wonder why there is any need to defend the belief-based conception against the swamping problem.

only requires that knowledge is *sometimes* more valuable than mere true belief, which is enough to generate an inconsistency when (S1), (S2), and (S3) are taken together.

(S2) maintains that once a belief is true, no additional instrumental value can make the belief more valuable. To see this point, consider an analogy by Linda Zagzebski (2003). Suppose that there are two cups of coffee and that one is brewed by a premium coffee machine, and the other by a terrible machine. Both cups of coffee taste equally good. If we value good coffee, we naturally value the good coffee machine rather than the bad coffee machine. However, if a cup of coffee tastes good, it does not matter whether it is brewed by a good machine or a bad machine, since how the coffee is brewed is only instrumental to the taste of the coffee. Similarly, if true belief is the only final value, how the belief is produced is only instrumentally valuable relative to the value of true belief. Once the belief is true, it no longer matters how the belief is produced. Accordingly, no additional instrumental value can make a true belief more valuable.

Given (S1) and (S2), it is difficult to see how knowledge can be more valuable than mere true belief. Given (S2), the property that confers additional value on true belief must be finally valuable. Given (S1), however, no property of belief, except being true, is finally valuable. Therefore, given (S1) and (S2), there is no way that knowledge can be more valuable than mere true belief. (S1), (S2), and (S3) are inconsistent.

To solve the swamping problem, one must reject at least one of the theses. Most philosophers take (S2) for granted (Sosa 2003, 161-63; Pritchard 2010, 16),[34] and I think they are right. Some philosophers (Sartwell 1992) reject (S3). However, since (S3) only requires that knowledge is *sometimes* more valuable than *mere* true belief, (S3) looks plausible. Therefore, the only available option is to reject (S1).

To reject (S1), one must find the Meno property of belief, *if* the belief-based conception of epistemic evaluation is adopted. The reason is that the belief-based conception assumes that the value of knowledge consists solely of the values of properties of belief. Hence the belief-based conception imposes a constraint on the answer to the swamping problem: the value of knowledge must be explained in terms of some valuable properties of belief.

[34] The only exceptions I found that reject (S2) are Adam Carter and Benjamin Jarvis (2012) and Adam Carter et al. (2013). I will compare my solution to the swamping problem with theirs in Section 6.

In the next section, I will discuss two theories of knowledge, *simple reliabilism* and *the virtue-theoretic account*. It is often assumed that the virtue-theoretic account could solve the swamping problem, whereas simple reliabilism suffers from the swamping problem. However, contrary to general opinion, I use Zagzebski's coffee analogy to show that, if the belief-based conception is adopted, the virtue-theoretic account does no better than simple reliabilism in accounting for the value of knowledge. However, if the belief-based conception is replaced with the performance-based conception, we can easily see, in Section 4, that both simple reliabilism and the virtue-theoretic account can offer the same solution to the swamping problem.

3.3 The Meno Property and the Belief-Based Conception

In this section, I will show how the belief-based conception of epistemic evaluation renders the swamping problem insurmountable by examining the virtue-theoretic account's objection to simple reliabilism, on the grounds that simple reliabilism suffers from the swamping problem. Zagzebski's coffee analogy is often used to show that simple reliabilism is affected by the swamping problem. However, I will show that the coffee analogy works *only because of the belief-based conception*. Also, if the belief-based conception is adopted, the virtue-theoretic account will suffer with the swamping problem just as simple reliabilism does.

Many proponents of the virtue-theoretic account (Riggs 2002; Sosa 2003; Zagzebski 2003) use the swamping problem to undermine *simple reliabilism*, which holds the view that knowledge is true belief produced through a reliable belief-formation process. The argument against simple reliabilism is roughly as follows: for simple reliabilism, the only difference between knowledge and mere true belief is *the property of being reliably-produced*. Therefore, *being reliably-produced* must be the Meno property of belief. However, reliability is usually defined in terms of *truth-conduciveness*: a belief-forming process is reliable if and only if it is likely that the beliefs it produces are true. Therefore, a reliably-produced belief is more likely to be true. But it seems absurd to say that a more-likely-to-be-true true belief is more valuable *qua* belief than mere true belief. Since the former belief is already true, adding the property of being-more-likely-to-be-true does not confer additional value on true belief. To put it differently, it is plausible that being reliably-produced is valuable for being more-likely-to-be-true, but it is valuable only instrumentally or as a means to the value of truth. Once the belief is true, it no longer matters whether the belief is reliably produced. Therefore, the property of being

reliably-produced cannot be the Meno property. Simple reliabilism cannot solve the swamping problem.

Zagzebski's coffee analogy can be used to show this point. Compare two cups of coffee, one brewed by a good coffee machine and the other by a bad coffee machine. Presumably, being brewed by a good coffee machine could make a coffee more likely to be tasty. However, once we taste both cups of coffee and find them equally tasty, adding the property of being-brewed-by-a-good-coffee-machine (being-more-likely-to-be-tasty) will not make the coffee brewed by the good coffee machine *better qua coffee* than the cup of coffee brewed by the bad coffee machine. Similarly, once the belief is true, the property of being-reliably-produced will not make a true belief *better qua belief*.

Proponents of the virtue-theoretic account argue that it can solve the swamping problem (Zagzebski 1996; 2003; Riggs 2002; Sosa 2007; Greco 2010). There are some substantial differences among them, especially between Zagzebski and the rest. Zagzebski's version of the virtue-theoretic account is often called *virtue responsibilism*, and the versions by John Greco, Wayne Riggs and Ernest Sosa are called *virtue reliabilism* (Battaly 2008). Zagzebski even argues that the virtue reliabilists also suffer the swamping problem. Unfortunately, although Zagzebski's objection to simple reliabilism is well-received, her objection to virtue reliabilism seems to be ignored in the literature. The reason, as I will argue in this chapter, is that Zagzebski gives a wrong diagnosis of (and a wrong solution to) the swamping problem. I will return to Zagzebski's view in Section 6. In this section, my aim is to use Zagzebski's coffee analogy to identify the cause of the swamping problem. I will use John Greco's theory as the representative version of the virtue-theoretic account, because it provides a nice contrast to simple reliabilism. According to Greco (2010, 12), S knows that p, if and only if:

(i) 'p is true';

(ii) 'S believes that p'; and

(iii) 'S believes the truth because S's belief is produced by intellectual ability'.

An intellectual ability, according to Greco (2010, 10), is reliable, so his virtue-theoretic account is also a version of reliabilism (see also Greco 1999; 2010). The only difference between Greco's virtue-theoretic account and simple reliabilism is (iii). (iii) expresses the core idea of the virtue-theoretic account: the fact that the knower believes the truth that p is explained in terms of the knower's exercise of her cognitive abilities rather than sheer luck. Accordingly,

while simple reliabilism holds that knowledge is a true belief produced by an intellectual ability, Greco maintains that knowledge is a true belief whose chief cause is the exercise of an intellectual ability.

To illustrate their difference more clearly, consider the well-known Barn Façade case (Goldman 1976). Suppose Barney sees and forms a true belief that there is a barn in front of him. Unknown to Barney, however, he is walking in the 'barn façade county', where all barns are fake, except the one that Barney sees is real. According to simple reliabilism, Barney *knows* that there is a barn in front of him (suppose Barney's vision is reliable). According to Greco (2010, 76-80), however, Barney does not know that there is a barn in front of him, because the most salient cause of Barney's true belief is not the exercise of his visual perception, but sheer luck.

Since simple reliabilism is criticised on the grounds that the property of being reliably-produced cannot be the Meno property that increases the value of true belief *qua* belief, it is natural to ask what the Meno property is, according to Greco's virtue-theoretic account. It seems that the only candidate available to Greco is the property of being a-true-belief-whose–chief-cause-is-the-exercise-of-an-intellectual-ability. Or, to put it simply, *the property of being true-because-reliably-produced*. As Greco answers the Meno problem, 'virtuously produced true belief is more valuable than both true belief that is not virtuous and virtuous belief that is not true' (2010, 99). So the virtue-theoretic account clearly considers being true-because-reliably-produced can make a true belief more valuable *qua* belief.

But how does the property of being true-because-reliably-produced confer additional value on true belief? Consider the coffee analogy again. The analogy invites us to evaluate two cups of coffee. Both taste equally good, and the only difference between them is that one is brewed by the reliable coffee machine and the other by the unreliable coffee machine. Let us compare the virtue-theoretic account with simple reliabilism:

> Simple Reliabilism
>
> GOOD: the coffee is tasty *and* brewed by the reliable coffee machine;
>
> BAD: the coffee is tasty *and* brewed by the unreliable coffee machine.
>
> The Virtue-Theoretic Account
>
> GOOD: the coffee is good *because* brewed by the reliable coffee machine;

BAD: the coffee is good and brewed by the unreliable coffee machine.[35]

As the coffee analogy reveals, our intuition is that simple reliabilism would render the values of coffee in both cases as the same. The explanation is simple: the reason why the property of being-reliably-brewed is valuable is that a reliably-brewed coffee is more likely to taste good. But this means that the property of being-reliably-brewed is only instrumentally valuable relative to the value of tasty coffee. Since the coffee is tasty, the property of being reliably-brewed cannot confer additional value to the coffee.

Let us examine whether the virtue-theoretic account fares any better. The virtue-theoretic account must be able to explain why the coffee in the good case is more valuable than the one in the bad case. The only explanation available to the virtue-theoretic account is that the coffee in the good case is *tasty because brewed by the reliable coffee machine*. So the property of being *tasty-because-reliably-brewed* must be able to increase the value of tasty coffee. Moreover, since the virtue-theoretic account argues that the property of being reliably-brewed does not increase the value of tasty coffee, it must explain how a tasty-*because*-reliably-brewed coffee is more valuable than a tasty-*and*-reliably-brewed coffee.

It is difficult to see how a tasty-because-reliably-brewed coffee is more valuable than a tasty-and-reliably-brewed coffee. Note that in the coffee analogy, the only difference between the two cups of coffee is that one is brewed by a good coffee machine and the other by a bad machine. Hence the chief cause of the good taste of the first cup of coffee seems to be the fact that it is brewed by a reliable coffee machine, while the chief cause of the good taste of the second cup is sheer luck. However, our intuition about the analogy is that the two cups of coffee are equally good *qua* coffee. Our reason is that what really matters about coffee is the taste. If the coffee tastes good, it does not matter how the cups of coffee are brewed (in the case of simple reliabilism), nor does it matter whether they are tasty because reliably brewed (in the case of the virtue-theoretic account).

The virtue-theoretic account needs to explain why being tasty-because-reliably-brewed is finally valuable. But it is difficult to see how it is so. According to the virtue-theoretic account, to say that the coffee is tasty because reliably brewed is to identify the way the coffee is brewed as the *chief* cause of its good taste among other causes, such as the quality of beans,

[35] I deliberately formulate the bad cases as the same for the sake of my argument. The point I want to highlight is whether the property of being good-because-brewed-by-the-reliable-machine adds to the value of good coffee.

water, etc. However, since the way the coffee is brewed is only instrumentally valuable relative to the value of coffee, it is difficult to see why identifying a cause of merely instrumental value as the chief cause can increase the final value of the consequence. The fact that being brewed by the reliable machine is the chief cause does not change the fact that being brewed by the reliable machine is merely instrumentally valuable to the value of tasty coffee. Hence, I think that the virtue-theoretic account cannot explain how being tasty-because-reliably-brewed can confer additional value to a cup of tasty coffee. Similarly, the property of being true-because-reliably-produced also cannot confer final value to true beliefs. Hence, the virtue-theoretic account does no better than simple reliabilism with respect to finding the Meno property.

If the coffee analogy can be used to show that simple reliabilism suffers the swamping problem, I have argued that it can show that the virtue-theoretic account also faces the swamping problem. The assumption behind the coffee analogy is that what is under evaluation in each case is *solely* coffee and belief. One cannot just answer that a coffee brewed by a reliable machine is better because a reliable coffee machine is also finally valuable. This is because what is under evaluation is not the coffee machine itself, but the coffee brewed by the machine. The question is whether the fact that the coffee is brewed by a reliable machine makes a difference to the value of coffee *qua* coffee. Similarly, in the case of epistemic evaluation, the coffee analogy assumes that what is under evaluation is solely *belief*. Hence, simple reliabilism and the virtue-theoretic account cannot just say that a reliable process or a virtuous exercise of one's ability is finally valuable, because the question is whether being reliably-produced or being true-because-reliably-produced can make a difference in the value of true belief *qua* belief.

Indeed, Zagzebski does not deny that a reliable process or ability is finally valuable. What Zagzebski wants to argue is that reliability, or more generally, the cause of a belief, cannot confer its value on true belief. As Zagzebski explains:

> What the [coffee] analogy shows is not only that a reliable cause does not confer value on its effect but also that there is a general problem in attributing value to an effect because of its causes, even if the value of the cause is independent of the value of the effect. I am not suggesting that a cause can never confer value on its effect. Sometimes cause and effect have an internal connection, such as that between motive and act, which I shall discuss in a moment. My point is just that

> the value of a cause does not transfer to its effect automatically, and certainly not on the model of an effect as the output of the cause. So even if the cause of true belief has an independent value, that still does not tell us what makes knowledge better than true belief if knowledge is true belief that is good in some way other than its truth. The second moral of the value problem, then, is this: *Truth plus an independently valuable source cannot explain the value of knowledge.* (Zagzebski 2003, 14)

Zagzebski (2003, 13-14) intends to use the coffee analogy to refute simple reliabilism as well as the virtue-theoretic account by Greco and Sosa. Unfortunately, Zagzebski's insight does not seem to be fully appreciated. For example, the entry 'The Value of Knowledge' in *The Stanford Encyclopedia of Philosophy* has the following remark on Zagzebski's objection based on the coffee analogy:

> There is reason to think that this objection will only at best have an impact on [simple] *reliabilist* proposals—i.e., those views which treat all reliable belief-forming processes as conferring a positive epistemic standing on the beliefs so formed. For example, the kind of [virtue] *reliabilist* account offered by John Greco might be thought to be untouched by this sort of argument. This is because, according to [virtue] reliabilism, it is not any sort of reliable process which confers positive epistemic status to belief, but only those processes that are stable features of what Greco calls the agent's "cognitive character". … Plausibly, however, one might argue that the reliable traits that make up an agent's cognitive character have some value independently of the instrumental value they possess in virtue of being reliable—i.e., that they have some final or intrinsic value. If this is right, then this opens up the possibility that agent-reliabilists can evade the value problem that Zagzebski identifies for pure reliabilists. (Pritchard and Turri 2014)

Note that this remark does not explain how an independent valuable reliable trait confers its value on true belief, while it does indicate that simple reliabilism cannot explain how a reliable process confers its value to true belief. Its assumption is that simple reliabilism cannot explain why a reliable process is finally valuable, but the virtue-theoretic account can show that a reliable trait is finally valuable. The idea is that T-monism is the cause of the swamping problem.

However, not just any rejection of T-monism is able to solve the swamping problem. As we have seen, Zagzebski accepts that reliable traits could be finally valuable (so she also rejects T-monism). What she rejects is the claim that those traits can confer their value on true beliefs (if the beliefs are considered as separate products). The critical assumption of the coffee analogy, therefore, is the belief-based conception. What the belief-based conception does to the swamping problem is to impose a constraint: epistemologists must explain why knowledge is *a more valuable kind of belief qua belief* than mere true belief. Accordingly, finding an independently valuable cause of true belief is not enough to save the value of knowledge from being swamped.

Zagzebski thinks that the value of knowledge will be swamped if one assumes that belief is a separate product of a cognitive process. I think that she is wrong. I agree with the above remark that the value of reliable traits can explain why knowledge is more valuable than mere true belief. However, that would mean that the critical move to solve the swamping problem is to reject the belief-based conception. In the next section, I will argue the reason why the virtue-theoretic account can solve the swamping problem is that it adopts the performance-based conception instead of the belief-based conception.

3.4 Achievement and the Performance-Based Conception

I have argued that the virtue-theoretic account also fails to account for the Meno property. However, this is not the way the proponent of the virtue-theoretic account presents its solution to the swamping problem. Greco does not claim that he can solve the swamping problem by arguing that being true-because-reliably-produced can make a true belief more valuable *qua* belief. Instead, Greco argues that knowledge is a kind of *success because of intellectual ability*, that is, a kind of *achievement*. A mere true belief would be just a lucky success, and an achievement would be more valuable for its own sake than a lucky success.[36]

I think that Greco's solution to the Meno problem is basically right. As I have argued, however, if one adopts the belief-based conception, one has difficulty in explaining how knowledge is a more valuable belief *qua* belief than mere true belief. So I think that the reason why the virtue-

[36] Duncan Pritchard (2010) and Daniel Whiting (2012b) discuss that some achievements are not valuable at all. But since solving the swamping problem only requires that knowledge is *sometimes* more valuable than mere true belief, Greco only needs the thesis that achievement is sometimes more valuable than lucky success. I critically discuss the virtue-theoretic account of achievement is more detail in Chapter 4.

theoretic account can successfully solve the Meno problem is that it implicitly rejects the belief-based conception. The value of an achievement, as I will explain in more details later, consists of two distinct final values: *the value of success* and *the value of excellent performance*. In the case of knowledge, its value consists of *the value of true belief* and *the value of virtuous cognitive performance*. My account of the value of knowledge is thus directly contrary to the belief-based conception of epistemic evaluation, which maintains that the value of knowledge must consist solely of the values of belief. Instead, I adopt *the performance-based conception of epistemic evaluation*, which holds that the value of knowledge can consist of some value lying outside belief (even granted that knowledge is a kind of belief,). For the performance-based conception, therefore, to solve the swamping problem needs not to find the Meno property. It offers a simple solution to the Meno property: knowledge is more valuable than mere true belief because knowledge has the value of virtuous cognitive performance but mere true belief does not.[37] My view is in direct opposition to Zagzebski's: my performance-based conception holds that truth plus an independently valuable source *can* explain the value of knowledge.

I will explain the performance-based conception more in the next section. What I want to do now is to show that although the virtue-theoretic account often appears as though it still works under the belief-based conception, it is better interpreted as adopting the performance-based conception. Let's first see how Greco explains why the virtue-theoretic account can solve the swamping problem:

> We have been working with the idea that knowledge is a kind of success from ability. … Put yet another way, in cases of knowledge *S* gets things right *because* she is intellectually able and because she has exercised her abilities. But now an answer to the value problem falls out of this account straightaway. In the *Nicomachean Ethics* Aristotle makes a distinction between (a) achieving some end by luck or accident, and (b) achieving the end through the exercise of one's abilities (or virtues). It is only the latter kind of action, Aristotle argues, that is both intrinsically valuable and constitutive of human flourishing. "Human good," he writes, "turns out to be activity of soul exhibiting excellence." … If this is correct then there is a clear difference in value between knowledge and mere

[37] This solution cannot solve the more demanding value problem of knowledge: how knowledge is more valuable than gettierised belief. I discuss this value problem in Chapter 4.

> true belief. In cases of knowledge, we achieve the truth through the exercise of our own intellectual abilities, which are a kind of intellectual virtue. (Greco 2010, 97-98)

Greco maintains that knowledge is a belief being true because the knower is intellectually able and she exercises her intellectual abilities. Greco thus argues that knowledge is more valuable than mere true belief, because exercising intellectual abilities is constitutive of human flourishing.[38] But that means that what really explains the extra value of knowledge over true belief is not any property of belief, but the exercise of intellectual ability. Put differently, when someone does not know but only has true belief—like the guide in Socrates' story—his lack of knowledge could be due to the fact that his intellectual ability is defective or he does not exercise his ability properly, which means that he is a less intellectually flourishing person.[39] That is why knowledge is more valuable than mere true belief. When someone knows, she not only has a true belief, which is valuable for its own sake, but also manifests her intellectual excellence, which is also valuable for its own sake.

The thesis that knowledge is an achievement shows that the proponent of the virtue-theoretic account actually adopts or ought to adopt the performance-based conception of epistemic evaluation. Someone's achievement is a success that she achieves through her ability. So when we evaluate whether someone's performance is her achievement, we evaluate whether her performance is *successful*, whether her performance is *competent* (whether she is able and she exercises her ability properly), and whether the success is *attributed* to her competence.[40] Knowledge, according to the virtue-theoretic account, is 'a kind of *achievement*, or a kind of success for which the *knower deserves credit* [my italics]' (Greco 2010, 7). Hence, epistemic evaluation is not just about evaluating the value of belief. It also evaluates the knower: how intellectually competent she is.

---

[38] I present my own account of the idea that exercising intellectual abilities is constitutive of human flourishing in Chapter 5.

[39] His lack of knowledge can be due to the fact that his belief is gettierised. One can have gettierised beliefs, even if the beliefs are true and one has those beliefs through exercising one's intellectual abilities. In that case, gettierised beliefs can still contribute to one's intellectual flourishing. I discuss the value problem about gettierised belief in Chapter 4.

[40] Here I explain epistemic evaluation in the structure of Ernest Sosa's account (Sosa 2007; 2011). Both Greco and Sosa hold that knowledge is more valuable than gettierised belief, because a true-because-competent belief is better than a true-and-competent belief. In Chapter 4, however, I argue that their view is wrong. My view is that the epistemic value of knowledge consists entirely of epistemic success (true belief) and epistemic competence.

In other words, how the virtue-theoretic account solves the swamping problem is not to show that there is a property of belief that confers final value on true belief. Instead, it rejects the belief-based conception of epistemic evaluation and opts for the performance-based conception. The virtue-theoretic account, therefore, could hold that the value of knowledge partially consists in the value of good cognitive performance, which is finally valuable.

Thus, to solve the swamping problem, the virtue-theoretic account rejects T-monism by maintaining that good cognitive performance itself is also finally valuable. The crucial point here is that it is difficult to reject T-monism if the belief-based conception is adopted.[41] If epistemic evaluation is all about the value of belief, it is difficult to see which property can make a true belief more valuable *qua* belief. As I have argued above, being reliably-produced or being true-because-reliably-produced appear to have only instrumental value to true belief. For belief's own sake, what matters is only whether it is true, or it does not really matter how it is produced. Hence, although it is true that the swamping problem is caused by T-monism, what makes T-monism look so plausible is the belief-based conception. As Zagzebski's coffee analogy makes the point vivid, it is difficult to see how an independently valuable source can confer its value on true beliefs. To solve the swamping problem, therefore, I suggest that we reject the belief-based conception.

I have suggested that the virtue-theoretic account is better interpreted as adopting the performance-based conception instead of the belief-based conception. However, simple reliabilists can solve the swamping problem by adopting the performance-based conception: simple reliabilists can also say that knowledge is more valuable than mere true belief because a successful exercise of the reliable ability is finally valuable. As Wayne Riggs points out, there are two assumptions behind the criticisms of simple reliabilism from the virtue-theoretic account:

> The critics' arguments rest on two false assumptions about reliabilism. They assume, first, that reliabilism need be committed to the view that the entire value of knowledge is constituted by the value of the known belief (hereafter, 'the belief assumption'). Second, they assume that reliabilism must be interpreted

[41] Could one adopt the performance-based conception and remains a T-monist? There might be such a possibility. But T-monism would lose its appeal if the performance-based conception is adopted. It is difficult to see why true belief is the only epistemic final value. Under the performance-based conception, a good cognitive performance would also be epistemically finally valuable.

> along instrumentalist lines (hereafter, 'the instrumentalist assumption'). That is, they assume that the *only* value involving reliable processes is the instrumental value of their propensity to produce true, rather than false, beliefs. (Riggs 2002, 81)

Riggs maintains that the swamping problem is solved, if either the belief assumption or the instrumentalist assumption is rejected. However, as I discussed above, if the belief-based conception ('the belief assumption', in Riggs's terms) is not rejected, it is difficult to see how the property of being reliably-produced can add final value to true beliefs. That is, it is the belief-based conception that makes the instrumentalist assumption plausible. On the other hand, once the belief-based conception is rejected, it is easy to see that a proper exercise of a reliable cognitive ability is finally valuable. It is unfortunate to miss that the critical move in the virtue-theoretic account is the rejection of the belief-based conception rather than the rejection of the instrumentalist assumption. In the literature on the swamping problem, many (e.g. Pritchard 2010, ch.1; Carter and Jarvis 2012; Sylvan 2012) wrongly assume that the solution to the swamping problem is to reject the instrumentalist assumption. It is the performance-based conception of epistemic evaluation, I think, that is the real contribution of the virtue-theoretic account to the philosophical reflection on the value of knowledge.

I have argued that the belief-based conception is the cause of the swamping problem and the performance-based conception offers a straightforward account of the value of knowledge. However, one may question whether the performance-based conception is correct. After all, knowledge is a kind of belief (or a mental state) rather than performance. In the next section, I will offer some defence and a more detailed account of the performance-based conception.

3.5 The Performance-Based Conception of Epistemic Evaluation

The performance-based conception regards epistemic evaluation as the same as performance evaluation. In this section, I will explain and defend the performance-based conception in more detail. (I offer a more comprehensive account of the performance-based conception in Chapter 2.)

When we evaluate a performance, it necessarily involves the performance *per se* and the *result* of the performance, as Sosa puts:

> Performances relate in one direction to *the agents involved*, and in another direction to *their performance-distinct products*, if any. So they might be evaluated with a view towards one direction, or towards the other. … Performances are in this way double-faced. So the evaluation of a performance seems not after all independent of the evaluation of the other two components of the performance situation. The evaluation is either *agent-involving* or *outcome-involving* (or both). Creditable performances must be attributable to the agent's skills and virtues, and thus attributable to the agent himself. (Sosa 2003, 167; my italics)

There are two aspects of performance evaluation: the result of the performance and the performance itself (or the performer). In the case of cognitive performance, a cognitive performance is the exercise of a cognitive ability in order to form a true belief. Epistemic evaluation can thus be divided into two components: evaluating the truth of the belief and evaluating the competence of performance.[42] I shall discuss the value of each component in turn.

Performance evaluation is to evaluate a performance *qua* performance. A performance usually aims at a certain result. A performance can be evaluated as *successful* if and only if it achieves the result. Moreover, a performance can also be evaluated as *competent* if and only if it is skilful (i.e., it exercises the ability and skill properly). Hence, a performance is good *qua* performance if it is successful or competent. Success and competence can enhance the value of performance *qua* performance.

Consider the coffee analogy again. Coffee brewing aims for good coffee. It is good *qua* coffee brewing if the resulting coffee tastes good or it brews the coffee skilfully. Furthermore, we can also evaluate each component of the performance for its own sake: a cup of coffee is good *qua* coffee if it tastes good; (roasted) coffee beans are good *qua* coffee beans if the coffee brewed from them tastes good; a coffee machine is good *qua* coffee machine if it can brew good coffee reliably. We can also give a more comprehensive account of skilful coffee brewing: whether the coffee beans are ground properly, whether the coffee machine is operated

[42] Sosa thinks that being apt—being successful because competence—also makes a performance more valuable. However, I argue that aptness does not make a performance more valuable in Chapter 4. For the present purpose, we can safely ignore this issue, since the concern here is whether knowledge is more valuable than mere true belief. The fact that a competent performance is finally valuable is sufficient for knowledge to be more valuable than true belief.

correctly, whether the temperature of water is suitable for coffee brewing, whether the brewing time is correct.

Undoubtedly, the reason why we value good coffee brewing is mainly that we value good coffee. However, that does not mean that the value of good coffee brewing is entirely instrumental. True, coffee brewing is instrumental in having a cup of coffee. However, the value of being good *qua* instrument or means is not derived from the value of the end (Kagan 1998; see also Chapter1, section 2). A Nazi's gas chamber has its own value of being good *qua* gas chamber even if its end—massive killing—is extremely evil. The value of the Nazi's gas chamber *qua* gas chamber consists in its capability to fulfil the function of massive killing rather than the value of massive killing. Therefore, the value of the Nazi's gas chamber as an effective instrument of massive killing is not merely instrumental relative to the value of massive killing, but finally valuable as being good *qua* gas chamber or *qua* instrument of massive killing. Similarly, the value of a good coffee machine is finally valuable.

Let's turn to cognitive performance. A cognitive performance aims at true belief. A cognitive performance is good *qua* cognitive performance if and only if the belief it produces is true or the performance is intellectually competent or skilful. Similarly, we can evaluate each component of cognitive performance for its own sake: a belief is good *qua* belief if it is true; an intellectual ability is good *qua* intellectual ability if it is reliable in producing or maintaining (certain types of) true beliefs (each kind of intellectual ability, such as perception or reasoning, can also be evaluated for its own sake). Furthermore, we can also evaluate whether a cognitive ability is exercised well. For example, Frege had a terrific ability of logical reasoning, but he might occasionally commit himself to some logical paradoxes. That does not show that his capacity for logical reasoning was bad, just that some exercises of his logical reasoning are bad. So when a cognitive performance results in knowledge, there is a true belief, which is finally valuable, as well as a reliable ability that produces the true belief properly, which is also finally valuable. Therefore, knowledge is more valuable than mere true belief.

One may object that the performance-based conception of epistemic evaluation is wrong, because knowledge is not a performance but rather a belief or a mental state.[43] However, it is wrong to assume that a state cannot be evaluated *qua* performance. To respond to this objection, the distinction between the performance itself and the result of performance is

[43] I address this objection in more detail in Chapter 1.

crucial. We often evaluate a performance through evaluating the result of the performance. The reason is usually that the performance was completed and all we have now for evaluation is the result of the performance. For example, people can go to see Michelangelo's statue of David in Florence. Although the performance of Michelangelo's sculpting David was done five hundred years ago, we can evaluate his performance through evaluating the statue of David. We can praise Michelangelo for achieving success, because the sculpture is so vivid and beautiful. We can also praise the skilfulness of Michelangelo's performance, because his skills are manifest in the sculpture. So we can praise the statue of David as successful and skilful. But when doing this, what we are really praising is Michelangelo's performance as successful and skilful. So even if the statue of David is not the performance itself, it can still be evaluated in terms of performance evaluation and its value consists of the value of a successful statue and the value of a skilful performance. Similarly, belief is a result of cognitive performance. We can also evaluate a belief as successful and competent, but what is really successful and competent is the cognitive performance.

So we can evaluate (in a derivative sense) a state *qua* performance. But why do we, when evaluating whether a belief is knowledge, evaluate a belief *qua* performance rather than *qua* belief? The reason is that knowledge is an achievement, or more precisely, a state that one can be in only through proper cognitive performance. So to evaluate whether one's belief is knowledge, one must evaluate not only whether the belief is true but also whether one gets the true belief through proper cognitive performance.

I have argued that to solve the swamping problem, we need to replace the belief-based conception of epistemic evaluation with the performance-based conception. One may object that the belief-based conception can also solve the swamping problem. In the next section, I will consider different proposals maintaining that the belief-based conception can solve the swamping problem.

3.6 Comparison

In this section, I will compare my solution to the swamping problem with several recent proposed solutions. They still maintain the belief-based conception. I will argue that my solution is better than theirs, and that their insights can be captured by my solution.

3.6.1 Linda Zagzebski

We have seen that Zagzebski uses the coffee analogy to criticise simple reliabilism (as well as Greco's and Sosa's virtue-theoretic account). Her reason is that simple reliabilism assumes the *machine-product model of knowledge*: knowledge is an 'external product of a believer in the way that the cup of espresso is the external product of the machine' (2003, 15). Zagzebski then argues that, as the coffee analogy shows, whatever theory assumes the machine-product model falls victim to the swamping problem. She suggests that the machine-product model should be replaced by the view that a belief is an act taken by an agent:

> If we think of a belief as part of the agent, the belief can get evaluative properties from features of the agent in the same way that acts get evaluative properties from the agent…. An act is not a product of an agent but is a part of the agent, and the agent gets credit or discredit for an act because of features of the agent. In particular, an agent gets credit for certain good features of an act, for example, its good consequences or the fact that it follows a moral principle – because of features of the act that derive directly from the agent – for example, its intention or its motive. If believing is like acting, we have a model for the way the agent can get credit for the truth of a belief because of features of the belief that derive from the agent. I propose, then, that this is the third moral of the value problem: *Knowing is related to the knower not as product to machine but as act to agent*. (Zagzebski 2003, 15-16)

Zagzebski claims that, since belief is an action, a belief can acquire valuable properties from the motives of an agent, just as the value of an act can be affected by the agent's motives. Accordingly, Zagzebski argues that knowledge is an intellectually virtuous true belief, and one component of an intellectual virtue is the love of truth, which is intrinsically valuable. Zagzebski (2003, 17-18) thus argues that knowledge is more valuable than mere true belief.

Zagzebski's diagnosis of the swamping problem is problematic in many ways. First, the idea that a belief is literally an act done by an agent is controversial. I do not want to claim that this idea is impossible to defend (though I argue that belief is a state in Chapter 2). For my purposes here, it suffices to say that my proposal does not require this controversial idea. Second, even if some instances of belief can be regarded as motivated acts, it is unlikely that all beliefs are motivated acts. A sudden burst of thunder makes me know that lightning just occurred, though I do not have a motive to know that. A similar problem is that not all

knowledge is motivated by the love of truth. A wife may desperately disregard all evidence that her husband has had an affair. However, she may one day unexpectedly see her husband having sex with his lover. Or a father attends the murder trial concerning the death of his daughter only because he wants to hear that the defendant is sentenced as guilty. Given that the evidence presented in court conclusively shows that the defendant really is the murderer of his daughter, the father knows who the murderer is. However, the father does not attend to the trial out of wanting to know the truth, but only out of wanting revenge. It is implausible to deny knowledge in those cases. Zagzebski would deny, I think, that knowledge in those cases is more valuable than truth, because they are not motivated by the love of truth. Zagzebski at best explains why *some* knowledge is more epistemically valuable than mere true beliefs. But my proposal shows how *all* knowledge is epistemically more valuable.

Moreover, as Jason Baehr (2012) argues, the love of truth may not be always intrinsically valuable. Suppose that someone particularly loves to know about the private lives of other people. He loves to read gossip columns. He reads his friends' emails and phone messages. He certainly knows lots about people's private lives. But his motive is hardly valuable. Hence, knowledge from such motivation would not become more valuable, if not becoming less valuable.

So, I think that Zagzebski's account of the value of knowledge is not successful. Let me revisit her reason for the claim that a belief is an act that is part of an agent. As we have seen, she concludes from the coffee analogy that '*truth plus an independently valuable source cannot explain the value of knowledge*' (2003, 14). Since the source cannot be separate from the belief if we want to explain the value of knowledge, Zagzebski then suggests that to explain that we need to regard beliefs as acts of the agent.

The conclusions that Zagzebski and I draw from the coffee analogy are in direct opposition. My solution is exactly that 'truth plus an independently valuable source' can explain the value of knowledge. The reason why Zagzebski cannot see my solution as an alternative is that she still maintains the belief-based conception. She is right that given the belief-based conception, an independently valuable source cannot confer additional value to true belief. Hence, she adopts the controversial thesis that belief is an action instead, because that thesis expands the conception of belief to include motivation as part of the belief. But I think that Zagzebski makes a mistake here. What the coffee analogy shows is not that our conception of knowledge as a kind of belief is wrong, but that our conception of epistemic evaluation is wrong. The fact

that knowledge is a kind of belief does not mean that when we evaluate the value of knowledge, its value consists solely of the values of belief. The reason is that when we evaluate whether a belief is knowledge, our concern is not just about belief but also about the process that brings about the belief. Under the performance-based conception, we can still maintain the 'machine-product model of knowledge' that Zagzebski criticises. As I discussed above, the fact that the performance and the result of the performance are separate does not make performance evaluation inapplicable to the result of the performance. When my friend buys a new coffee machine and gladly invites us to taste the coffee it brews. We find that the coffee tastes good, so we judge the machine as being good. We may ask my friend which model it is and where to buy it. Certainly, there is no problem to evaluate the value of coffee machine through evaluating the coffee it brews. Whether they are separate does not matter. Hence, the performance-based conception does not require us to reject the machine-product model. Zagzebski's reason for the idea that belief is an act is not adequate.

3.6.2 J. Adam Carter, Benjamin Jarvis, and Katherine Rubin/ Miranda Fricker

J. Adam Carter, Benjamin Jarvis, and Katherine Rubin (2013; see also Carter and Jarvis 2012) try to solve the swamping problem by rejecting (S2). They argue that an instrumental value can add extra value to knowledge. For the coffee analogy, which lends support to (S2), is a disanalogy. They invite us to consider the example of a pleasant home. Although the property of being well-maintained is merely instrumental to a pleasant home, a well-maintained pleasant home is still more valuable than a pleasant home. Hence, instrumental values are not always swamped.

The reason is that maintaining a home, unlike brewing a cup of coffee, does not have a precise terminus.[44] Carter et al. think that instrumental value is only swamped when the process of producing the relevant final value is over, as in the case of coffee brewing. When the machines are still brewing coffee, and supposing that both cups of coffee are good, they argue that we should choose the coffee brewed by the good machine instead of the one by the bad machine. The instrumental value of being brewed by a good machine is not swamped by the tastiness of the coffee. It suggests that so long as the project of maintaining a house is not over, the instrumental value of being well-maintained cannot be swamped by the pleasantness of the

[44] Carter et al. attribute their idea to Matthew Chrisman (2012) who discusses the evaluation of maintaining beliefs, which is discussed in Chapter 1.

home. Similarly, since belief is a kind of state, being in a state of true belief often requires good management. The property of being well-managed—that is, being justified—can confer additional value, despite being instrumentally valuable, on true belief. The swamping problem is thus solved.

Their solution captures Socrates's insight that knowledge is more stable than mere true beliefs. However, their explanation is not right. It is not at all clear that instrumental value can confer additional value when the terminus has not been reached. Consider the coffee analogy again. Carter et al. seem to think that the fact that we should choose the cup of coffee brewed by the good machine indicates that the instrumental value of a good machine is not swamped so long as it is still brewing. However, this demonstrates their confusion between ontological and epistemological matters. The reason why we should choose the cup of coffee brewed by the good machine is simply the fact that since we do not know how each cup of coffee will taste, it is wiser to choose the coffee brewed by the good machine if we want a good cup of coffee. The key factor is not whether the machine is still brewing the coffee. Rather, it is the fact that we do not know what the coffee tastes like. That fact does not show that we think that a reliably-brewed good coffee is better than a merely good coffee. The instrumental value has already been swamped, though we just do not know it yet.

In the case of a well-maintained pleasant home, it is intuitive that we do value it over a merely pleasant home. However, Julien Dutant (2013) argues that the example does not show that an instrumental value of being well maintained is not swamped by the already-present final value of a pleasant home. Instead, it is swamped by the home being pleasant *today*, but not by the home being pleasant *tomorrow*. Since we want our home to be pleasant tomorrow, we certainly want our home to be well maintained. But that does not solve the swamping problem, because the problem is about how knowledge is more valuable than mere true belief *now*.

This may not be a serious problem to Carter et al. if Miranda Fricker (2009) is correct to argue that the swamping problem rests on an unwarranted *synchronic presumption*:

> According to which the value question is conceived as a question about the comparative values of mere true belief and knowledge at a snapshot in time. At best, we are invited to compare a mere true belief that p and knowledge that p in a very short time frame. (Fricker 2009, 127)

Fricker also mentions Socrates's insight that knowledge is tethered and argues from it that Socrates rejects the synchronic presumption. For Fricker, knowledge is more valuable than mere true belief, because it is more resilient—more likely to survive misleading counter-evidence—than mere true beliefs.[45] Therefore, Fricker and Carter et al. may be happy to accept the fact that knowledge is no more valuable than mere true beliefs when evaluated synchronically, but will reject that it should be taken as a problem.

However, my solution can capture Socrates's insight without giving up the thesis that knowledge is *always* more valuable than mere true beliefs. According to my proposal, when one knows, it means that one has a good cognitive ability and one exercises the ability well. The fact that one can exercise a good cognitive ability well implies that one's known true belief is more resilient. However, the extra value of knowledge is explained in terms of the value of being good *qua* cognitive ability and being good *qua* exercise of cognitive ability. My solution to the swamping problem remains workable even under the synchronic presumption. It shows that my solution is preferable to those of Fricker and Carter et al.

3.6.3 Berit Brogaard

Berit Brogaard (2006) argues that the rejection of the belief-based conception on the ground of its inability to solve the Meno problem is based on the false assumption that things of the same intrinsic properties are equally valuable. Brogaard accurately observes that the core insight of the virtue-theoretic account is the replacement of the belief-based conception of epistemic evaluation: 'The focus shifts from evaluating the belief itself to the state the agent is in when she is responsible for believing something true. … When the extra credit the knower is due does not make the known belief more valuable, it supposedly adds value to the overall state of knowing that p' (Brogaard 2006, 336). However, she argues that in order to solve the Meno problem, it is not the case that the belief-based conception must be discarded. Brogaard (2006, 339-41) argues that the reason why the virtue-theoretic account holds that the property of being reliably-produced cannot confer value on true belief is that the property of being reliably produced is an *extrinsic property*. Furthermore, the virtue-theoretic account assumes a Moorean conception of value, according to which things that have the same *intrinsic properties* are equally valuable. Accordingly, a reliably-produced true belief is equally

[45] However, Jonathan Kvanvig (2003) and John Hyman (2010) argue that true beliefs are often as resilient as knowledge.

valuable as a mere true belief. However, Brogaard points out, citing Rabinowicz and Rønnow-Rasmussen (2000), that extrinsic properties can add final value. For example, Princess Diana's dress is more valuable than an exact copy, even though the property of belonging-to-Princess-Diana is an extrinsic property of the dress. So Brogaard argues that the virtue-theoretic account is wrong to rule out the possibility that the property of being reliably produced may confer value on beliefs.

However, I think that Brogaard's criticism is mistaken. The objection of the virtue-theoretic account to simple reliabilism is not based on the Moorean conception of value, because the property of being-true is arguably an extrinsic property of most beliefs (Berker 2013a, 344 n.11). If proponents of the virtue-theoretic account accepts the Moorean conception, they would also agree that being true does not make beliefs more valuable. Obviously, this is not the case, because the swamping argument also assumes that being true is finally valuable and being reliably-produced is instrumentally valuable relative to the value of true belief.

Instead, I think that the objection to simple reliabilism does not rest on any controversial idea. As the coffee analogy shows, the value of being brewed by a reliable coffee machine is entirely instrumental relative to the value of good coffee. Once the cup of coffee tastes good, it does not matter whether it is brewed by a reliable machine. Although it is true that an extrinsic property may add value, Brogaard does not really show that being reliably-produced can add value to a cup of coffee or to a true belief. For there is an important disanalogy between Princess Diana's dress and a reliably-produced belief. People value Princess Diana's dress, *because* they value Princess Diana. If they do not value Princess Diana, it would be very puzzling why they value her dress. However, it is not the case that we value reliably-produced belief, *because* we value a reliable belief-formation process. It is more likely the case that we value reliably-produced belief, because the belief is more likely to be true. It shows that being reliably produced is entirely instrumentally valuable relative to true belief. Brogaard's diagnosis of the swamping problem, therefore, is wrong.

3.7 Conclusion

I have argued that the belief-based conception is the cause of the swamping problem. To properly explain the value of knowledge, we should adopt the performance-based conception. The performance-based conception is compatible with the orthodox view that knowledge is a kind of belief (or a mental state): it holds that knowledge is a kind of belief that one can only

have through proper cognitive performance. The performance-based conception offers a simple solution to the Meno problem: knowledge is more valuable than mere true belief, because knowledge has the extra value of good cognitive performance. That is, the reason why the guide who knows the road to Larissa is better than the guide who merely has the true belief is that the former is intellectually more competent than the latter. We value knowledge, not only because we value true belief, but also because we value people being intellectually competent. This is a point missed by the belief-based conception.

Recently, epistemologists became interested in an updated version of the Meno problem: why is knowledge more valuable than gettierised belief? Virtue epistemologists often claim that the virtue-theoretic account can solve this problem. In the next chapter, however, I will argue that once the performance-based conception is properly understood, virtue epistemologists should conclude that knowledge is no more valuable than gettierised belief. What it shows is that the purpose of epistemic evaluation is not really about knowledge, but about epistemic success and competence.

Chapter 4.

Why Value Knowledge? Virtue Epistemology and the Gettierian Value Problem

4.1 Introduction

In the last chapter, I have dealt with Socrates' Meno problem: Why is knowledge more valuable than mere true belief? The problem is central to philosophical reflection on the nature of knowledge, especially the element(s) that distinguishes knowledge from mere true belief. Recently, the problem about the value of knowledge has been recast as what I will call the *Gettierian value problem* (Jones 1997; Kvanvig 2003; Pritchard 2010): Why is knowledge more epistemically valuable than mere justified true belief?[46] The Gettierian value problem arises because of the Gettier problem (Gettier 1963). The Gettier problem shows that justified true belief is not sufficient for knowledge. An answer to Socrates' Meno problem can be that some components of knowledge—e.g. justification—make knowledge more valuable than mere true beliefs. But that cannot answer the Gettierian value problem. The Gettierian value problem is more demanding than the Meno problem.

There is another reason why contemporary epistemologists are interested in the Gettierian value problem. After decades of failure in solving the Gettier problem, epistemologists have started to reflect on the importance of the Gettier problem. It is in this context that epistemologists have become interested in the Gettierian value problem. If knowledge is no more valuable than gettierised belief, perhaps it is not important whether the Gettier problem

[46] Note that the Gettierian value problem, unlike the Meno problem, is about only the epistemic value of knowledge. This is how most epistemologists understand the Gettierian value problem and is my concern in this chapter. I am not concerned with the prudential or non-epistemic values of knowledge.

is solvable, and epistemologists should turn to more important issues (Kaplan 1985; Jones 1997; Kvanvig 2003).

Virtue epistemology (Greco 2010; 2011; Sosa 2007; 2011) claims that the virtue-theoretic account of knowledge can satisfactorily answer the Gettierian value problem. The virtue-theoretic account of knowledge has a unique feature that distinguishes it from other theories of knowledge: that is, the virtue-theoretic account considers knowledge as a kind of cognitive performance and thus epistemic normativity as a species of performance normativity, as discussed in Chapter 2. We evaluate a cognitive performance in the same way that we evaluate a musical or athletic performance. Ernest Sosa explains that there are three aspects of performance evaluation: whether the performance is *successful*; whether the performance is *competent* or *adroit*; and whether a performance is *apt*, that is, whether it is successful because competent. Similarly, according to the virtue-theoretic account, a cognitive performance is successful when the belief is true, competent when the performance is done through exercising one's cognitive abilities or virtues, and apt when the performance is successful because competent. When apt, on this account, the belief is knowledge.[47]

Virtue epistemologists claim that this account provides an adequate answer to the Gettierian value problem. The Gettierian value problem, rephrased in terms of the virtue-theoretic account, is why an apt performance is better *qua performance* than a successful, adroit but inapt performance.[48] The core argument of the virtue-theoretic account is that a performance that is successful *because* competent (apt) is *a better performance* than one that is successful

[47] In Chapter 2, I have argued that although belief can be assessed in terms of performance evaluation, what really falls under performance evaluation is cognitive performance rather than belief. This is an important difference between my view and Sosa's. When I say that a belief is competent, I mean that the belief is produced by a competent performance.

[48] Note that the phrase '*qua* performance' is important here. Since the virtue-theoretic account regards knowledge as a kind of performance, the value of knowledge that the virtue-theoretic account explains is the value of knowledge *qua performance*, i.e. the value of being an apt performance. The Gettierian value problem is generally regarded as a problem about the *epistemic value* of knowledge. One may question whether the virtue-theoretic account can really answer the Gettierian value problem on the grounds that the value *qua* performance may not be an epistemic value. Sosa's account (2007, ch.4) of epistemic value is that it is a kind of value within *the epistemic domain*, and belief and knowledge, as cognitive performances, fall within it. According to Sosa, therefore, the value of knowledge *qua* performance is the epistemic value of knowledge. Although I accept the virtue-theoretic account of epistemic evaluation as a kind of performance evaluation, I reject Sosa's account of epistemic value. I propose my account of epistemic value in Chapter 1 and discuss the problem of Sosa's account in Chapter 5. However, that difference does not affect anything in this chapter. In the next section, it can be seen more clearly that what the virtue-theoretic account explains is an account of the value of knowledge *qua* performance.

because of luck (inapt). Since performance evaluation applies to belief and knowledge, it seems that the virtue-theoretic account offers a neat solution to the Gettierian value problem.

Unfortunately, although the virtue-theoretic account of knowledge and its value is elegant and simple, I will argue in this chapter that its answer to the Gettierian value problem is wrong. The reason is that virtue epistemologists do not fully grasp the thought that epistemic evaluation is a kind of performance evaluation. In the next two sections, I will first argue that the virtue epistemologists' argument fails to prove that knowledge is more valuable than gettierised beliefs, and then I shall argue that what the virtue-theoretic account actually shows, contrary to the intent of virtue epistemologists, is that knowledge is no more valuable than gettierised beliefs.

A caveat. In this chapter, my aim is a modest one. I do not argue that knowledge *is* as valuable as gettierised belief. I only argue that *if* one accepts the virtue-theoretic account, one should regard knowledge as no more valuable than gettierised belief. Neverthelss, if my account is correct, it can offer a solution to one of the major objections to the virtue-theoretic account, namely, that it cannot solve the Gettier problem. I discuss how the virtue-theoretic account can solve the Gettier problem in the appendix of this chapter.

4.2 The Virtue-Theoretic Account of the Gettierian Value Problem

According to the virtue-theoretic account, knowledge is more valuable than gettierised beliefs, because a performance that is successful *because* competent is better *qua* performance than one that is successful because of luck, as John Greco and Sosa explain:

> Suppose [a] belief to be competently acquired but Gettiered, so that it is true only by epistemic luck. … Inapt performances fall short not only in that they might have been *better* on relevant dimensions. They fall short in the fuller sense that they fail to meet minimum standards for performances. Because they are inapt, they are therefore *flawed*: not just improvable but defective. (Sosa 2011, 46)
>
> The proposed solution answers even [the Gettierian value problem], respecting the supposition that knowledge is more valuable than all of its parts taken together. This is because success from ability is more valuable than an act that is both successful and from ability, but not successful *because* from ability. Suppose, for example, that an athlete runs a race in a way that is clearly an exercise of her

> athletic excellence. Suppose also that she wins, but only because the other runners, some of whom are equally excellent, get sick before the race. Or suppose that she wins, but only because the other runners were bribed. Clearly, neither sort of win is as valuable as it could be. What one really values as an athlete is to win as the result of ability. (Greco 2010, 99)

Greco and Sosa both hold that a successful-because-competent performance is better than a successful-because-lucky performance. But their reasons are different. Sosa argues that since being inapt fails to meet the standard for performances, inapt performances are, to that extent, flawed and defective *qua* performance. However, Sosa's account simply assumes that since aptness is one aspect of performance evaluation, aptness makes a performance better *qua* performance. Greco, on the other hand, offers us an example to demonstrate why apt performances are better than inapt ones. I will argue in section 3 that aptness itself does not make a performance better *qua* performance. In this section, my aim is to show that Greco's example fails to show that aptness makes a performance better *qua* performance.

Greco's athlete example seems to demonstrate that apt performances are better than inapt ones. Intuitively, the win because of the athlete's own ability (hence, an apt performance) ranks above those win because of others' physical difficulties (luck) or even cheating (hence inapt performances). Since performance evaluation applies to belief and knowledge, it seems that the virtue-theoretic account offers a neat solution to the Gettierian value problem.

However, our intuition about those examples of inapt performances cannot *universally* support the proposed solution offered by the virtue-theoretic account. True, when we say that a performance is successful because lucky, we tend to think that it is inferior to a performance that is successful because competent. However, that is because when people hear, 'He wins because of luck', people tend to think that *he would not win without the luck*. And from the claim that the winner would not win without the luck people tend to infer that the winner is not competent enough. That is, in conversation, people tend to infer from 'He wins the game because of luck' that he is *not competent enough* to win the game. It is not surprising, therefore, that we agree with virtue epistemologists that a performance that is successful because competent is better than a performance that is successful because lucky, since we tend to think that a successful-because-luck performance is *less* or *not competent*. This tendency is actually manifest in the quotation from Greco. Greco first correctly stipulates being inapt as '*not* "successful because adroit"'. He then moves from that stipulation to claim that

the athlete's inapt performance is 'successful *only* because lucky' (similarly, in Sosa's quotation, a gettierised belief is 'true *only* by epistemic luck'). Obviously, this is invalid. 'A performance is successful *only* because lucky' entails 'the performance would not be successful without the luck', which suggests that the performance is not competent enough.[49] But inapt performances are just 'not "successful because adroit"'. They may still be successful even without the luck. They could be as successful and competent as apt performances. Clearly, the virtue-theoretic account does not meet the Gettierian value problem head-on. The Gettierian value problem is whether apt performances are more valuable than inapt (gettierised) performances that are *equally successful and competent*. To avoid unnecessary confusion, the virtue-theoretic account needs to explicitly stipulate the examples of gettierised performance being inapt but *equally successful and competent*. Once this stipulation is explicitly made, however, it is unclear that we do have the intuition that an apt performance is better *qua* performance than an inapt performance that is equally successful and competent. Let us redescribe Greco's example and test our intuition:

> Suppose, for example, that an athlete runs a race in a way that is clearly an exercise of her athletic excellence. Suppose also that she wins, because the other runners, some of whom are equally excellent, get sick before the race. Or suppose that she wins, because the other runners were bribed. *But remember: she, in fact, is truly competent. That is, even if other athletes were in good conditions and competed in the race competently, she might still win the game*.

It is unclear to me that the athlete's performance is no as good simply because of being inapt in such a way. Indeed, it is difficult to see why we should rank such an inapt performance below the apt performance. In both cases, her performances were equally successful and competent. The only difference is that in the inapt case, it was just that other athletes happened to be ill or bribed (not even by her), so that it is hard to attribute the victory to the

[49] In light of this, it is difficult to see how Greco can describe the athlete's performance as successful only because lucky, but meanwhile maintains that her performance is competent ('an exercise of her athletic excellence'). For if she does run competently, it seems unfair to say that she wins *only* because other runners are sick or bribed. Saying that suggests that she would simply not be competent enough to win the race, if other runners were really competing. Hence, it would be wrong to describe it as a case in which the performance is competent. Competence does not just require that the performer does the best she can. Rather, it requires that the performer can reliably bring success (Greco 2010, 77). If the athlete can win the race *only* because other runners do not run competently, it is difficult to say that her running is competent. Therefore, Greco does not provide a convincing example that an apt performance is better than an inapt performance that is successful and competent.

winner's performance.[50] But it seems unfair that we thereby devalue her performance only because of an external factor beyond her control. Indeed, as I will argue later, luck does not affect the value of performance. A performance that is successful and competent is as good as a performance that is successful because competent.

I have argued that the virtue-theoretic account's answer to the Gettierian value problem is wrong. It simply misses the target. So far we have only tested our intuition about the value of being apt or inapt by considering one example. In the next section, I will look more closely into the notion of aptness. I argue that once the notion of aptness is properly understood, we should be able to see that apt performances are not necessarily better than inapt ones.

4.3 The Value of Being Apt: On the Gettierian Value Problem

The only difference between knowledge and gettierised beliefs is aptness. We have seen that aptness is usually explained in terms of 'success because of competence'. The phrase 'success *because of* competence' indicates that the point of evaluating the aptness of performance is to offer a causal explanation of the success. More specifically, aptness evaluation is not just seeking any cause of the success. It is to assess whether the performance's being successful can be explained by the performance's being competent. When that is the case, the success is *attributable* to the performance for its competence, as Greco explains:

> In cases of knowledge, S believes the truth *because* S's belief is produced by ability. In Ernest Sosa's terminology, S's belief is 'true *because* competent.' But what are phrases of the form 'X because Y' getting at in this context? The answer is attribution: In cases of knowledge, S's success is *attributable* to S's ability, which is to say that it is attributable to S. (Greco 2012, 1; see also Greco 2010)

Aptness evaluation, in other words, is a matter of identifying whether the competence of the performance is the *most* salient cause of the success. Why the most salient cause, not *the* cause of success? The reason is that for every event, there are usually multiple causes contributing to it, and not all of them contribute equally. If lightning causes a forest fire, lightning is not the only possible cause. Other causes include the dry and hot weather, oxygen,

[50] Some may object that it is unclear that the athlete's performance, when redescribed in this way, is still inapt. Since she may win anyway, it might be better to judge her performance as apt. However, the fact that it could be difficult to judge whether a performance is apt is an important lesson that we can learn from virtue epistemology but unfortunately remains ignored. I will discuss this point more later.

combustible materials such as woods and grasses. Without any one of these factors, the forest fire might not occur. But usually we think that lightning is the most salient cause, so we attribute the cause of the forest fire to lightning. More accurately, the fact that the forest fire occurs is *primarily attributable* to lightning, since there are other less salient causes.[51] In ordinary speech, nonetheless, saying that 'X is attributable to Y' usually conveys the idea that 'X is primarily attributable to Y'. Henceforth, when I say that 'X is attributable to Y', I meant that Y is the most salient, but may not be the only, cause of X.[52]

Now we see that when virtue epistemologists claim that a performance is apt, they mean that the success of the performance is attributable to the competence of the performance. In the light of this account of aptness, I would like to argue that the following thesis is true:

> (1) An apt performance is no more valuable *qua* performance than an inapt performance, if they are equally successful and competent.

(1) is contrary to what virtue epistemologists claim. More importantly, what I attempt to argue is that if aptness is understood as being a matter of attributing the success of a performance to the competence of the performance, as virtue epistemologists explain, virtue epistemologists should reject the assumption that lies behind the Gettierian value problem, i.e., that knowledge is more valuable than gettierised beliefs.

But is (1) true? It involves three properties of a performance: being successful, being competent, and being apt. According to Sosa's account, these properties are three essential standards of performance evaluation. It seems that, on this account, each of them is a *goodness standard for performance*; that is, satisfying each standard makes a performance good *qua* performance, as Sosa clearly holds in the quotation in the previous section. Although

---

[51] Greco (2010; 2012) and Sosa (2007; 2010) modify their accounts of aptness or attributability. However, their accounts are criticised by J. Adam Carter (forthcoming) and Lisa Miracchi (2015). I adopt Carter's account of aptness (forthcoming) that a performance is apt if and only if its success depends on the performance's being competent more than on luck. Accordingly, the success of the performance is primarily attributable to the performance rather than luck.

[52] This leads to a thorny problem: How do we know which cause is the most salient one? One might argue that if it was raining heavily, the forest fire would not happen even though many woods were struck by lightning. That is true, but we need not worry about this problem here. Indeed, the fact that there could be several causes that compete as the most salient cause is one of the most significant points about the virtue-theoretic account that has not been fully appreciated. For the present, we need to bear in mind only that, when virtue epistemologists claim that a performance is apt, they mean that the success of the performance is attributable to the adroitness of the performance

I believe that success and competence are the goodness standards for performance, *aptness is not*. Therefore, I hold that the following theses are true:

(2) A performance is good *qua* performance if it is successful;

(3) A performance is good *qua* performance if it is competent;

(4) Being apt is not a goodness standard for performance.

(2) and (3) are intuitive enough. It seems obvious that we judge successful performances as being better than unsuccessful performances *qua* performances and competent performances as better than incompetent performances *qua* performances. We seem to judge, for example, that if a chess player won a game, then his performance was better *qua* chess performance than if he lost it. But even if he lost the game, we would still judge that his performance was good *qua* chess performance if he played it skilfully. Hence, I grant that (2) and (3) are truisms. The problem is, of course, (4). Why is (4) true? And if it is true, why is aptness one crucial aspect of performance evaluation? I will answer the last question in the next section. Let me now answer the first question. Consider the following two cases of football game:

APT-FOOTBALL. England was fighting against Argentina. Both played hard and well. Both had many chances to score a goal, but failed. Finally. England scored a goal, and defeated Argentina 1-0.

INAPT-FOOTBALL. England was fighting against Argentina. Both played hard and well. In the first half, England scored a goal, but was disallowed because of an offside call. Unfortunately, the call was misjudged. Near the end of the game, England scored another goal. However, this time it actually was offside, but the referee did not notice it. Luckily, England won the game.[53]

Should we judge that England's performance in INAPT-FOOTBALL was not as good as England's performance in APT-FOOTBALL? One cannot just reply, as virtue epistemologists do, that a performance that is successful because competent is better than a performance that is successful because lucky. For this reply, as argued above, often suggests that the performance in INAPT-FOOTBALL is less competent. To highlight that the issue is about the value of being apt, let us stipulate the cases that England's performances in both cases were equally

[53] This is an example of gettierised performances, as it is constructed in accordance with Zagzebski's famous recipe for Gettier-style examples (Zagzebski 1994).

competent. Keeping this stipulation in mind, we should recognise that England's performances in both cases are equally good *qua* performances. The reason is that, had England's goal not been disallowed in INAPT-FOOTBALL, England would have won the game anyway. The fact that England's goal was disallowed was sheer bad luck.[54] The misjudgement was irrelevant to England's performance. It is difficult to see why such bad luck should affect the value of England's performance. A moral we learn from this example is this:

> (5) A performance that is inapt only because of bad luck does not become less good *qua* performance.

By 'inapt only because of bad luck', I mean:

> (6) A performance is inapt only because of bad luck if and only if (i) it would be apt were the bad luck absent, and (ii) the presence or absence of the bad luck does not affect the degree of the success and the competence of the performance.

If (5) and (6) are true, a performance that is inapt only because of bad luck is as good as an apt performance that is equally successful and competent. It follows that being apt or inapt does not affect the value of a performance *qua* performance. Hence, (4) is true.

(6) is merely a stipulation of the notion of being inapt because of bad luck. It is to emphasise that the issue under discussion is whether an apt performance is more valuable than an inapt performance when they are equally successful and competent. Therefore, if virtue epistemologists want to defend their view, they have to reject (5).

One possible response from virtue epistemologists that we can find in the literature is that apt performances are a kind of *achievement* but inapt performances are not, and an achievement is more valuable than a performance that falls short of achievement.[55] As Greco explains:

---

[54] One may say that it was not a sheer bad luck, but injustice or unfairness. It is not clear how this difference affects our judgment on the value of England's performance. However, the judge did not deliberately treat England unfairly. He just made mistakes, and these kinds of mistake are allowed and often occurs in football games. This is a kind of bad luck that every team could suffer. I give a definition of bad luck below. This sort of 'unfair' judgment fits my definition of bad luck.

[55] In Chapter 2, I also endorse the idea that knowledge is a kind of achievement. However, my view is different from Greco's view. For Greco, only an apt performance is an achievement. However, I think that an achievement can be inapt, as long as it is competent. Hence, as I will explain shortly, I do not think that the idea that knowledge is a kind of achievement can show that knowledge is more valuable than gettierised belief.

> Knowledge is a kind of *achievement*, or a kind of success for which the knower deserves credit. And in general, success from ability (i.e. achievement) has special value and deserves a special sort of credit. This is a ubiquitous and perfectly familiar sort of normativity. Thus we credit people for their athletic achievements, for their artistic achievements, and for their moral achievements. We also credit people for their intellectual achievements. Epistemic normativity is an instance of a more general, familiar kind. (Greco 2010, 7)

Greco's achievement thesis consists of two claims: a definitional claim that achievement *is* apt performance, and a value judgment that achievement is valuable for being creditworthy.[56] These two claims are interconnected. Greco's reasoning seems to be this: we *credit* people for their achievement, so the achievement must be *creditable* to them; since inapt performances are not creditable to the performers, inapt performances are not achievements. So apt performances are more creditworthy and thus more valuable than inapt performances.

However, it seems wrong that inapt performances are not achievements. It seems to me that the England team in the INAPT-FOOTBALL achieves no less than it does in the APT-FOOTBALL. More importantly, even if only apt performances are achievements, Greco would be wrong in saying that apt performances are more creditworthy than inapt performances. The mistake is now recognised by Greco himself. In his recent paper (2012) I cited above, Greco recognises that the notion of creditability is ambiguous:

> [Virtue epistemologists] also use the phrases 'creditable to S's ability' and 'creditable to S'. But 'creditable' is ambiguous between 'praiseworthy' and 'attributable'. To avoid that confusion, I will use the latter term. (Greco 2012, 1 n.1)

It is clear that Greco now wants to avoid the implication that being apt is praiseworthy, so he explains aptness in terms of attributability rather than creditability. However, that means that we do not praise a performance simply because it is apt. Greco cannot hold that apt performances deserve 'a special sort of credit' that inapt performances do not. Otherwise, it is difficult to see why Greco wants to avoid the implication that a performance is praiseworthy simply because of being apt. Greco must conclude that being apt does not make a

[56] Duncan Pritchard (2010) and Daniel Whiting (2012b) also offer different objections to Greco's achievement thesis.

performance more praiseworthy. That is, an apt performance, other things being equal, does not deserve more credit than an inapt performance. Certainly, an apt performance is praiseworthy, but this is not because it is apt, but because it is competent. Aptness is a matter of attributability, and competence is a matter of praiseworthiness. This result should not be surprising, because we often praise people who do their best and demonstrates their great skills and talents but nonetheless fail to achieve success. We think that they still deserve credit even though they fail. Therefore, as I said earlier, I think that England in INAPT-FOOTBALL achieves no less than England in APT-FOOTBALL: in both matches they score a goal, and win the game. Since by stipulation in both matches they play equally well, they should deserve the same degree of credit. Hence, they are equally valuable.

The objection from achievement, therefore, is wrong on two fronts: first, it is wrong that an inapt performance cannot be an achievement; second, it is wrong to understand aptness in terms of praiseworthiness rather than attributability due to the ambiguity of creditability. Hence, an apt performance, other things being equal, is no more valuable *qua* performance than an inapt performance.

Contrary to the intent of virtue epistemologists, what the virtue-theoretic account shows is that knowledge is no more valuable than gettierised beliefs.[57] It should not be surprising, because if we take seriously the thesis that epistemic normativity is a kind of performance normativity, we often do not devalue a performance when some factors of luck are present. Under the virtue-theoretic account of performance normativity, *success* is about whether the performance successfully achieves its goal. *Competence* is about whether the skill or ability is good enough and whether the skill or ability is performed well. Each makes a performance better qua performance. *Aptness*, however, concerns merely whether the success is attributable to the competence of the performance. Luck undermines the value of performance (as such) only when it makes the performance unsuccessful or less competent.

[57] Other virtue epistemologists, such as Jonathan Kvanvig (2003) and Duncan Pritchard (2010), also argue that the virtue-theoretic account fails to show that knowledge is more valuable than gettierised belief, on the grounds that the virtue-theoretic account also falls victim to the Gettier problem. Hence, the virtue-theoretic account is not an adequate account of knowledge. In the appendix of this chapter, however, I will argue that the virtue-theoretic account can manage the Gettier problem. My point of this section is to show that, assuming that the virtue-theoretic account of knowledge is correct, what it shows is that knowledge is no more valuable than gettierised belief.

4.4 Why Being Apt? The Purpose of Epistemic Evaluation

I have argued that being apt does not add any value to a performance. If so, this raises an interesting question: What is the point of evaluating whether a performance is apt? It seems pointless to evaluate a feature of performance that adds no value to a performance. Strictly speaking, my account would imply that there is no need to evaluate whether a performance is apt. If we want to know whether a performance is good *qua* performance, all we need to evaluate is whether it is successful and competent.

However, I think that assessing whether a performance is apt is not pointless. As discussed, assessing whether a performance is apt is trying to offer a causal explanation of the success. More specifically, the point of assessing the aptness of a performance is to evaluate whether the performance itself is *competent enough for being the most salient cause of the success*. The reason is that a performance is apt, as I have repeatedly emphasised, if and only if the success of the performance is attributable to the performance for being competent enough, i.e. enough for being the most salient cause of the success. This is a conceptual truth about aptness: aptness is attributing success to competence as the most salient cause of the success. In other words, the point of aptness evaluation is to determine whether a performance is competent enough. Aptness itself does not add any new value to a performance.

Moreover, I think that it is also an empirical truth that people often evaluate whether a success is attributable to the performance itself as a means to assess whether the performance is competent. Consider archery as a kind of performance. Assessing whether a shot is successful is quite easy: we just need to see if the shot hits the target. However, assessing the competence of the shot is not so straightforward. The success of the shot is certainly evidence for the shot's being competent. But the shot may hit the target simply because of luck. To ascertain that the shot is really competent, we, therefore, often need to rule out other factors as significant causes of the success of the shot. That is, aptness evaluation is often used as a way to determine whether a performance is competent by ruling out other competing factors behind the success of the performance.

So it is not surprising that when England's fans say 'England beat Argentina because of their skilful play', they imply that England played sufficiently competently by claiming that their performance was apt. Similarly, when Argentina's fans say that 'England won just because of luck', they imply that England has not played competently and its play was not as good as

England's fans suggest. If so, it is now clear why the talk of a performance's being apt or inapt carries such implications, because the point of aptness evaluation is about evaluating whether a performance is competent enough. When a performance is apt, it is competent. But it does not follow, as I argued above, that when a performance is inapt, it is incompetent.

To recap, although being apt does not make a performance better *qua* performance, aptness evaluation is not pointless. Evaluating the aptness of performance is a useful way to determine whether a performance is competent enough for its success. In other words, the point of performance evaluation is success and competence. Aptness evaluation is just a means to evaluating success and competence. Hence, proponents of the virtue-theoretic account should hold that the point of epistemic evaluation is not really about whether one knows that p, but about whether one's belief that p is true and one is epistemically competent enough in holding the true belief that p. Accordingly, knowledge is no more evaluable than gettierised belief because gettierised belief is equally competent and true.

In view of this, one may be tempted to reject the virtue-theoretic account, on the grounds that it implies that knowledge is no more evaluable than gettierised belief. I think that this is wrong. First, some epistemologists (Kaplan 1985; Kvanvig 2003; Pritchard 2010) also maintain that knowledge is no more valuable than gettierised belief. Some others (Baehr 2009; Fricker 2009) argue that we do not have a pre-theoretical intuition that knowledge is more valuable than gettierised belief. So it is wrong to reject a theory of knowledge because it denies that knowledge is more valuable than gettierised belief.

More importantly, it may be a simple fact that the focus of epistemic evaluation is true belief and epistemic competence. The virtue-theoretic account of epistemic evaluation is actually in line with Edward Craig's genealogical account of knowledge:

> Human beings need true beliefs about their environment, beliefs that can serve to guide their actions to a successful outcome. That being so, they need sources of information that will lead them to believe truths. … On any issue, some informants will be better than others, more likely to supply a true belief. … So any community may be presumed to have an interest in evaluating sources of information; and in connection with that interest certain concepts will be in use.

> … To put it briefly and roughly, the concept of knowledge is used to flag approved sources of information. (Craig 1990, 11)[58]

Craig argues that the concept of knowledge is useful, because people in a society like ours have the need for good information. The concept of knowledge is used to flag good informants or good sources of information, namely, *truth* and *epistemic competence*. This is exactly what the virtue-theoretic account would hold, according to my interpretation. To say that a performance is apt is to say that a performance is competent enough in for its success. Hence, the concept of knowledge is a useful tool to flag truth and epistemic competence.

If all the concept of knowledge does is to flag truth and epistemic competence, one might ask: Why should we not discard the concept of knowledge and just have the concept of truth and competence instead? As I have said, aptness evaluation is just a means—but not the only means—to determine whether a performance is competent enough. So the concept of knowledge (or aptness) could be eventually removed from epistemic evaluation. I do not hold that necessarily we need the concept of knowledge.

Furthermore, the fact that we can flag truth and epistemic competence without the concept of knowledge does not show that Craig's genealogical account is inadequate. The reason is that the Gettier problem is a *philosophical* one. It is unlikely that when people first use the concept of knowledge, they aim not only to flag truth and epistemic competence, but also to flag which belief is not gettierised. Moreover, the Gettier problem is a fairly recent one. It seems true that laypeople use the concept of knowledge to flag competent and true beliefs all the while without recognising the possibility of such beliefs' being gettierised.[59] Finally, if the above discussion shows us anything, one would be that people tend not to take heed of the difference between a successful-and-competent performance and a successful-because-competent performance. People tend to say that the performance is successful because competent in order to stress the competence of the performance, whereas saying that the performance is successful because lucky often suggests that the performance is not competent enough. This feature shows that people often just use the concept of knowledge to flag truth and competence.

---

[58] Sosa (1991, 281-82), Pritchard (2010), and Greco (2012) also concur with this view. Christoph Kelp (2011) offers a different genealogical account, but his account still supports my view that the point of knowledge is to flag truth and competence.

[59] I discuss the Gettier problem in the appendix of this chapter, which may explain why people may not pay heed of the difference between knowledge and gettierised belief.

Recognising that the point of epistemic evaluation is about truth and epistemic competence, let me consider one last objection. One may argue that we do prefer our performance to be apt rather than inapt, so that it suggests that apt performances are more valuable than inapt performances. If you were one of the English players (or one of England's fans), would you prefer to be in APT-FOOTBALL or INAPT-FOOTBALL? I would certainly prefer my performance to be apt. It seems natural that people want that the success of their performance to be attributable to their own performance. However, does that not suggest that apt performances are better performances than inapt ones?

The answer is no. First, performance evaluation is evaluating whether a performance is good *qua* performance. Whether a performance is good *qua* performance is determined by the evaluative standards of performance rather than by our preferences.[60] For example, England's fans may prefer Argentina's players to play the game badly. However, the preference of England's fans does not make Argentina's performance bad *qua* performance if they do play the game well.

Furthermore, although I agree that we prefer our performances to be apt rather than inapt, this does not mean that a performance becomes a better performance simply because of being apt. The reason why we prefer our performances to be apt can be explained by the fact that we want our contribution to the success to be recognised. If our performance is judged as apt, our performance is recognised as competent enough. However, when our performance is inapt, that is, when the success is not attributed to us, the value of our performance is often unfairly overlooked and underappreciated. Indeed, the value of inapt performance is unfairly overlooked by virtue epistemologists. People, like virtue epistemologists, often fail to appreciate the subtle difference between *being inapt only because of bad luck* and *being successful because of luck*. The former conveys the idea that the performance is *equally* successful and competent, while the latter the idea that the performance is successful but *less* or even *not* competent. It is not surprising that people prefer their own performance to be apt, because they do not want their own performance to be wrongly regarded as incompetent. Imagine that in INAPT-FOOTBALL, some Argentina's fans might be complaining that England won only because of luck. England's fans understandably would feel annoyed. England's fans would be right to defend that England deserved the win because they had played well and it was merely a bad luck that their first goal was wrongly disallowed.

[60] I discuss the idea of goodness of kind in Chapter 1.

To avoid this sort of personal or interested-party preference, it is better to take a disinterested third-party perspective. From a disinterested third-party perspective, it should be clear that the value of a performance should remain intact if the performance is inapt merely because of bad luck. We can imagine that a Japanese football commentator would rightly remark that England's performances in both games were equally good. The commentator would explain that the aims of a football team are to play well and to win the game. In both games, England achieved both aims. He would also point out that the only difference in both games was the incompetence of the referees in INAPT-FOOTBALL and that should not affect the value of England's performance. Therefore, the fact that we prefer our performances to be apt does not show that apt performances are more valuable than inapt ones. It just reveals a human factor—that we want recognition. Knowing that, the point of evaluating aptness is that we can see how being evaluated as apt helps our performances in being recognised. But that does not mean that aptness by itself makes a performance better *qua* performance.

4.5 Conclusion

In conclusion, I have argued that virtue epistemologists' answer to the Gettierian value problem is wrong. They fail to fully appreciate the implication of their account of epistemic evaluation as a kind of performance evaluation. Virtue epistemologists should conclude that gettierised beliefs can be as valuable as knowledge. The reason for this is that success and competence are the sole goodness standards for performances. The point of epistemic evaluation, according to the virtue-theoretic account, is to assess whether a cognitive performance is successful (producing true beliefs) and competent. We do evaluate the aptness of the performance, and use it as a means to assess whether a performance is competent enough. If the virtue-theoretic account of epistemic evaluation is right, then it follows that the central questions in epistemology are not 'What is knowledge?' and 'Whether one knows?' Instead, epistemologists should be primarily concerned about 'What is an epistemically successful and competent person?' and 'How do we make one epistemically more successful and competent?' In other words, epistemology should not be a philosophy of knowledge, but a philosophy of good knowers (epistemic agents). For a philosophy of knowledge, even if knowledge is shown as no more valuable than gettierised belief, the Gettier problem remains central. For a philosophy of knowers, however, the Gettier problem is sidelined because two

persons are equally good knowers if they are epistemically equally successful and competent.[61] Virtue epistemologists have long called for a turn in epistemology from being belief-based to person-based (Battaly 2008). My virtue-theoretic approach is an answer to their call. The idea that epistemology should be person-based (or performance-based) is crucial. From Chapter 2 to the present chapter, I have shown that the focus of epistemic evaluation is not belief but rather epistemic performance and the cause of some value problems of knowledge are exactly the belief-based conception of epistemic evalution. In the next chapter, I will argue that to propely explain the normative force of epistemic norms, we should adopt a person-based conception of epistemic evaluation rather than a belief-based one.

Appendix. Virtue Epistemology and the Fake Barn Case

In the main text, I have assumed that the virtue-theoretic account of knowledge is correct and argue that the Gettier problem should be sidelined. However, one (Kvanvig 2003; Pritchard 2010) may argue that the virtue-theoretic account is wrong, because it also suffers the Gettier problem. Accordingly, one could argue that my account of the value of knowledge based on the virtue-theoretic account would also be problematic. In this appendix, I will argue that the virtue-theoretic account can take some Gettier-style counterexamples as borderline cases, so they should not be taken as counterexamples to the virtue-theoretic account. My argument is supported by some recent results made by experimental philosophers.

Here I will focus on Pritchard's objection to the virtue-theoretic account based on the fake barn case. Pritchard describes the case as follows:

> Barney forms a true belief that there is a barn in front of him by using his cognitive abilities. … Accordingly, we would naturally say that Barney's cognitive success is because of his cognitive ability, and so we would, therefore, attribute a cognitive achievement [an apt belief] to Barney. … The twist in the tail, however, is that, unbeknownst to Barney, he is in fact in 'barn façade county', where all the other apparent barns are fakes. Intuitively, he does not have knowledge in this case because it is simply a matter of luck that his belief is true. (Pritchard 2010, 35-36)

[61] I offer my account of good knower in Chapter 5.

Pritchard claims that since Barney's belief is apt and the fake barn case is widely regarded as a Gettier-style counterexample, the virtue-theoretic account of knowledge is wrong.

So Pritchard's objection from the fake barn case consists of two claims:

(i) Barney's belief that there is a barn is apt;

(ii) Barney's belief that there is a barn is not knowledge.

In response, Greco and Sosa offer two opposite replies. Sosa (2007) claims that Barney's performance is apt, but rejects most epistemologists' judgment that Barney does not know. Greco (2010; 2011) opts for the opposite response. He agrees that Barney does not know, but argues that that is because his performance is not apt. In other words, Sosa accepts (i) and rejects (ii), whereas Greco rejects (i) and accepts (ii).

Let us consider (i) first. As we have seen, the question about aptness is a question about whether the performance itself is the most salient cause of its success. Inevitably, there must be borderline cases for which it is hard to judge whether a performance is apt. The fake barn case seem to be a borderline case about the aptness of Barney's belief. Barney's perceptual capacity and his luck that he happens to see the real barn are both salient causes in explaining why Barney's belief is true. It would be difficult to judge whether his perceptual capacity is the most salient cause of his true belief. Why does Pritchard think that (i) is true? Pritchard argues that (i) is true because the fake barn case is structurally similar to the following example:

> FORCE FIELD: Archie again selects a target at random, skilfully fires at this target, and successfully hits it because of his skill. ... Suppose, however, that unbeknownst to Archie there is a force-field around each of the other targets such that, had he aimed at one of these, he would have missed. It is thus a matter of luck that he is successful, in the sense that he could very easily have not been successful. Notice, however, that luck of this sort does not seem to undermine the thesis that Archie's success is a genuine achievement. … That is, his success in this case is still primarily creditable to his archery abilities, even despite the luck involved in that success. (Pritchard 2010, 35)

Pritchard holds that Archie's shot is apt, for his shot is accurate because of his own archery skill. Once we accept that, Pritchard argues, we should conclude that Barney's belief is also apt, because both cases are structurally the same. The force field case, however, seems more

likely to be a borderline case about aptness. Pritchard claims that the shot remains Archie's achievement, but he also claims that 'it is thus a matter of luck that he is successful'. It seems that Pritchard commits himself to the claim—which I think correct—that Archie's skill and luck are both salient causes of Archie's success. But which cause is more salient? It seems that the artificiality of the case makes us unable to judge which cause is more salient. I think that it is wrong to hold that Archie's shot is manifestly apt. It is clear that the force field case is a borderline case of apt performance. Accordingly, we do not have solid evidence that (i) is true. Instead of affirming or denying that Barney's belief is apt, like Sosa and Greco do, we should conclude that it would be arbitrary to judge whether it is apt and (i) should be rejected. Hence, the fake barn case should not be taken for granted as a counterexample to the virtue-theoretic account.

How about (ii)? Sosa does reject (ii). This move puts Sosa in opposition to the mainstream view among epistemologists. The fake barn case is a textbook example of gettierised beliefs. Sosa tries to explain away epistemologists' intuition about the fake barn case by distinguishing between *animal knowledge* and *reflective knowledge*: roughly, animal knowledge is apt belief, and reflective knowledge is apt belief that the subject aptly believes to be apt. Sosa then explains that while Barney does not have reflective knowledge about the barn, his belief amounts to animal knowledge (Sosa 2007, ch.2; 2011).

However, it is difficult to see how Sosa's distinction between animal knowledge and reflective knowledge would persuade those epistemologists who hold that Barney does not know in the fake barn county. For they can accept Sosa's claim that Barney's belief is apt (true because of seeing the real barn), but still deny that Barney knows that it is a barn.

However, in light of the above discussion about aptness, I think that virtue epistemologists should in turn cast some doubts on epistemologists' intuition about the fake barn case. Since (i) cannot be clearly established as true, virtue epistemologists could then explain away epistemologists' intuition for the same reason that the fake barn case is a borderline case of knowledge. Since it is difficult to judge whether Barney's belief is apt, it means that according to the virtue-theoretic account it is difficult to judge whether Barney knows. It is thus reasonable that people may consider the fake barn case as a Gettier-style counterexample, although in fact we cannot conclusively establish that this is the case.

The idea that the fake barn case is a borderline case of knowledge can be supported by recent research by John Turri, Wesley Buckwalter, and Peter Blouw (2014; see also Colaço et al.

2014). Since it is widely agreed that one lesson from the Gettier problem is that knowledge is incompatible with luck, Turri et al. wanted to study how luck affects laypeople attributing knowledge. They designed seven types of scenario and asked the participants whether the subject knows. The detail of these type are described in Figure 1 (Turri et al. 2014).

Experiment 4: Description of stories used in the seven conditions of Experiment 4

Condition	Description
1. Knowledge Control	The stone Emma purchases is a diamond. She walks out of the store and nothing else happens.
2. Failed Threat	The stone Emma purchases is a diamond. A skilled jewel thief tries to steal it from her pocket before she leaves the store, but he fails.
3. Detection Similar Replacement	The stone Emma purchases is a diamond. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Someone secretly slips a diamond into Emma's pocket before she leaves the store.
4. Detection Dissimilar Replacement	The stone Emma purchases is a diamond. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Long ago, Emma's grandmother secretly sewed a diamond into the pocket of Emma's coat.
5. No Detection Similar Replacement	The stone Emma purchases is a fake. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Someone secretly slips a diamond into Emma's pocket before she leaves the store.
6. No Detection Dissimilar Replacement	The stone Emma purchases is a fake. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Long ago, Emma's grandmother secretly sewed a diamond into the pocket of Emma's coat.
7. Ignorance Control	The stone Emma purchases is a fake. She walks out of the store and nothing else happens.

Figure. 1 Seven Types of Knowledge Scenario

In all of these scenarios, after purchasing a stone, Emma forms a belief that the stone is a diamond. In all scenarios except Ignorance Control, Emma's belief that the stone is a diamond is true, but their causes are different. Two major factors figures in the explanation of the truth of Emma's belief: first, whether the stone that Emma perceives when purchasing it is a real diamond; second, whether the stone Emma purchases is replaced by another diamond and how the replacement happens. Excluding the second factor, the first factor alone determines two control cases: *Knowledge Control* (Emma knows that it is a diamond) and *Ignorance Control* (Emma falsely believes that it is a diamond). The second factor introduces different kinds of luck into the scenarios. The experiment is to compare those five scenarios with the two control cases to see how luck affects people's attribution of knowledge. The participants are asked whether Emma in each case knows. The idea is that if a rate of knowledge attribution in a case is sufficiently closer to Knowledge Control, it is more likely that people think that the agent in that case knows. The result (Turri et al. 2014) is as follows:

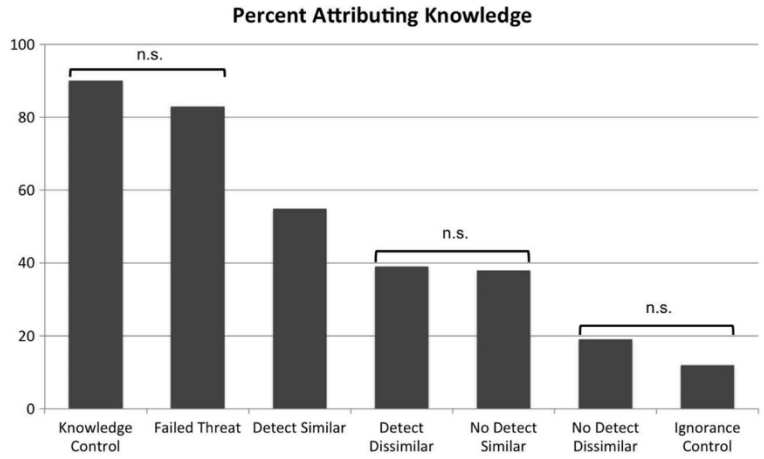

Experiment 4: Percentages of participants attributing knowledge across conditions. Except where nonsignificance is indicated, the significance for all comparisons is at the $p < .01$ level.

Figure. 2 Attribution of Knowledge in Each Scenario

Many Gettier-style examples can be classified as one of these types. Chisholm's sheep in the field case (Chisholm 1989) and Zagzebski's husband in the living room case (Zagzebski 1994) can be classified as No Detection Similar/Dissimilar Replacement. The experiment supports epistemologists' intuition that the beliefs in those cases are gettierised, because clearly they are much closer to Ignorance Control rather than Knowledge Control. The fake barn case, however, is an instance of Failed Threat. Barney does perceive the real barn. Although his belief is not safe, threatened by surrounding fake barns, the truth maker of his belief is never replaced. Interestingly, according to the experiment, Fail Threat is very close to Knowledge Control. It suggests that the fake barn case should not be considered as a Gettier-style counter-example, as Turri et al. conclude:

> Our findings suggest an explanation for the seemingly inconsistent prior findings and theorizing on Gettier cases: Knowledge attributions are sensitive to different forms and combinations of luck, and prior research on Gettier cases has not controlled for all of the sensitivities identified here. Indeed, … researchers would count the stories used for five of the separate conditions in Experiment 4 as Gettier cases: Failed Threat, Detect Similar, Detect Dissimilar, No Detect Similar, and No Detect Dissimilar. But if intuitions about Gettier cases vary as widely as our results indicate—from patterns that closely resemble responses to *paradigmatic knowledge* (Failed Threat) to patterns that closely resemble responses to *paradigmatic ignorance* (no detect dissimilar)—then 'Gettier case' is a theoretically useless category. The fact that something is a Gettier case would

> be consistent with its being both overwhelmingly judged knowledge and overwhelmingly judged ignorance, thereby masking differences that radically affect the psychology of knowledge attributions, and depriving the category of any diagnostic or predictive value. (Turri et al. 2014)

Experimental philosophy, however, is controversial, and many philosophers dispute its validity or its philosophical relevance.[62] I do not claim that we should reject the fake barn case as a Gettier-style counterexample simply because some experiments have been tried by experimental philosophers. But rather what we have here is a case when theory and empirical data mutually support each other. If knowledge is apt belief as the virtue-theoretic account suggests, it is not surprising that the experiment would have this result. Since aptness is a matter of attributing success to the performance over luck (or other non-performative factors), the virtue-theoretic account would predict that the more salient the luck is, the less likely that knowledge is attributed. The experiment confirms this predication. Since luck is less salient in Failed Threat than Detect Similar, it is expected, and in fact is the case, that the rate of knowledge attribution is higher in Failed Threat than Detect Similar, and so on.

In conclusion, in order to make the fake barn case a Gettier-style counterexample to the virtue-theoretic account, (i) and (ii) must be true. However, the virtue-theoretic account can reject (i), because the fake barn case appears to be a borderline case about aptness. The virtue-theoretic account can in turn explain why we do not have reliable intuition about whether the subject in the fake barn case knows. This explanation is supported by recent experiments by experimental philosophers, and the virtue-theoretic account offers a theoretical explanation why the experiments would have these results. Contrary to the view that the fake barn case is a counterexample to the virtue-theoretic account, the virtue-theoretic account instead puts our intuitions about Gettier-style cases into doubt.

[62] For a discussion, see Alexander and Weinberg (2007)

Chapter 5.

Epistemic Evaluation and Epistemic Normativity

> Any account that does explain the source of epistemic normativity must explain how it is that epistemic claims have normative force. If you tell me that a belief of mine is unjustified, this gives me reason to give up that belief. The epistemic claim is something about which I should care, and an account of the source of epistemic norms must explain why it is that I should care about such things.
>
> --Hilary Kornblith (2002, 145)

5.1 Introduction

In this chapter, I want to discuss the relationship between epistemic evaluation and epistemic normativity.[63] By epistemic normativity, I mean the normative force of epistemic norms or requirements. As Hilary Kornblith remarks above (and I will assume without arguments), the normative force of epistemic norms is *categorical* rather than *hypothetical*.[64] That is, there is an important difference between epistemic norms and, say, chess rules. Chess rules are

[63] One aim of this chapter is to address a problem I left in Chapter 1. In Chapter 1, I propose a new account of epistemic value in terms of goodness of epistemic kind in place of the orthodox view that epistemic value is a kind of intrinsic value from the epistemic point of view. In the end, I mentioned an objection to my proposal that goodness of epistemic kind cannot account for epistemic normativity. I aim to provide such an account in this chapter.

[64] William Alston (1988) famously argues against the possibility of epistemic norms on the grounds that we cannot voluntarily control our beliefs. I will not discuss his view here. For responses to Alston, see Philippe Chuard and Nicholas Southwood (2009), Richard Feldman (2000), Conor McHugh (2011b), and Sharon Ryan (2003).

hypothetical. The fundamental rule of chess, perhaps, is that one should checkmate. But *if* I do not want to play chess or I do not want to win the game, other things being equal, the rule that one should checkmate has no force for me at all. That means that the normative force of chess rules is conditional upon something other than the chess rule themselves. Epistemic norms, such as that one should not have unjustified beliefs, are not hypothetical. The fact that epistemic norms are categorical means that, *ceteris paribus*, one should not have unjustified beliefs, *period*. Even though I want to believe that global warming is a big fat lie, I should still not believe this, because such a belief is not justified. That does not mean that I should not have unjustified beliefs *all things considered*. The categoricity of epistemic norms is compatible with the fact that they can be overridden. Arguably, a father may unjustifiably believe that his son would return home before midnight, because his son asks the father to believe so and believing his son helps him maintain a good relationship with his son.

There are many epistemic norms, but in this chapter, I focus on two:

(F) One should not have false beliefs;

(J) One should not have unjustified beliefs.

We have seen that Kornblith accepts (J). (F) also seems plausible enough, because it expresses the idea that one of our epistemic goals is to avoid falsehood (James 1979; Lynch 2009a; Whiting 2012a). For the present purpose, I will assume that both norms are correct. The purpose of this chapter is to explain why (F) and (J) are categorical. More specifically, I want to examine *the evaluative approach* to epistemic normativity. By the evaluative approach, I mean the articulation of epistemic evaluative standards and the explanation of the normativity of epistemic norms in terms of the epistemic evaluative standards. My aim here is to examine whether the evaluative approach can account for epistemic normativity, that is, why epistemic requirements like (F) and (J) are categorical.

In particular, I will discuss two major evaluation approaches: *the value approach* and *the correctness approach*. The value approach takes epistemic value as its central notion and tries to explain epistemic normativity in terms of the epistemic values. The correctness approach, on the other hand, holds that epistemic normativity can be grounded in the correctness-conditions for belief.

The value approach usually tries to explain epistemic normativity in terms of the epistemic values of belief. The value approach has recently been criticised by Stephen Grimm (2009). I

will first discuss Grimm's argument and conclude that the value approach cannot explain epistemic normativity (section 2). I will then argue that the correctness approach also fails to explain epistemic normativity (section 3). I suggest that they are both unsuccessful, because the object under evaluation in both approaches is *belief*: the fact that a false or unjustified belief is epistemically disvaluable or incorrect is not normatively robust enough to generate categorical norms. To explain why epistemic norms are categorical, we need to shift from belief to *person*. I will develop a value approach based on the epistemic evaluation of person—call it *the personal value approach*—and explain how my account can account for epistemic normativity (section 4).

5.2 The Value Approach

Let me begin with a very rough account of the value approach. The value approach tries to account for epistemic normativity in terms of epistemic value. Most value approaches focus on the epistemic value *of belief*. Usually, true belief is considered the fundamental epistemic value and false belief the fundamental epistemic disvalue. Justification is often explained in terms of truth-conduciveness: justified belief is more likely to be true, and unjustified belief is more likely to be false. Hence, justification is also an epistemic value, despite not being fundamental.[65] Thus, the value approach attempts to explain the reason why we should, other things being equal, conform to epistemic norms, such as (F) and (J), because conforming to (F) and (J) helps promote epistemic values and reduce epistemic disvalues.

The question now is, 'What is *fundamental epistemic value*?' One way is to understand it as a kind of *intrinsic value*. This understanding is natural, because intrinsic value could give rise to categorical norms. The reason is that many philosophers (Bradley 2006; Brentano 1969; Broad 1930; Ewing 1948; Ross 1939) hold that if something is intrinsically valuable, the following claim is true:

> INTENTIONAL. If one evaluates x, without reference to circumstances or consequences, as intrinsically valuable, one must hold that, *ceteris paribus,* one should promote it, preserve it, or bring it about. On the contrary, if one evaluates

[65] There could be other versions of the value approach that do not take only true beliefs as fundamental epistemic values. The version under discussion is often called, more precisely, veritism or truth monism. Although I only discuss veritism, my own objection applies to non-veritistic versions of the value approach.

> x, without reference to circumstances or consequences, as intrinsically disvaluable, one must hold that, *ceteris paribus,* one should destroy it, prevent it, or stop it from happening.

INTENTIONAL is often used as a test for intrinsic value (as discussed in more details in Chapter 1). If happiness is an intrinsic value and pain is an intrinsic disvalue, then one should promote happiness and prevent pain. Since intrinsic value generates categorical norms, the value approach naturally would want to hold that epistemic value is a kind of intrinsic value. However, an immediate problem arises: not all true beliefs appear to be intrinsically valuable, and not all false beliefs are intrinsically disvaluable.[66] Some true beliefs are so trivial that it is difficult to see why we should promote or preserve them. Some true beliefs may even be intrinsically bad, such as true beliefs about the manufacture of massive killing weapons or about people's privacy, so we should prevent ourselves from believing them.[67] Similarly, some false beliefs are so trivial or insignificant that it is difficult to see why we should take effort to prevent or eliminate them. For example, after finishing reading *Being and Time*, I put it to the right of *Being and Nothingness*. However, I tend to confuse right and left, so I came to believe that I put *Being and Time* to the left of *Being and Nothingness*. This mistake is hardly significant. It is unlikely to make me unable to find the book in the future. And even if it would, it is important to note that INTENTIONAL asks us to consider the false belief by itself without consideration of the consequences or circumstances. It is difficult to see why such a false belief is intrinsically disvaluable.

One common way in which the value approach can deal with this problem is to weaken the epistemic value of true belief as intrinsically valuable *only from the epistemic point of view* (Pritchard 2014). Call it *the perspectival approach*. The phrase 'from the epistemic point of view' is widely used by epistemologists to differentiate between epistemic value and non-epistemic values. Limiting the value of true belief only to a certain perspective, the value approach could make it more defensible. Even so, there is strong resistance to the claim that all true beliefs are intrinsically valuable from the epistemic point of view, because trivial true beliefs still do not appear to be intrinsically valuable even from the epistemic point of view. Indeed, William Alston, who coined the phrase 'from the epistemic point of view', holds that

---

[66] I call this problem 'the triviality objection' in Chapter 1, where I discussed this objection more comprehensively. The point of discussing the triviality objection is to show how it poses a problem to the value approach to epistemic normativity.

[67] This point is made by Jason Baehr (2012).

only true beliefs about 'matters that are of interest or importance to us' (Alston 2005, 29; see also Goldman 1999; 2002) have epistemic intrinsic value. However, as Grimm (2009) argues, such a *restricted view*—according to which only some kinds of true (false) belief have epistemic intrinsic (dis)value— faces a serious difficulty. For if trivial beliefs have no intrinsic value, it follows that epistemic evaluation of these beliefs is out of place. A trivial true or false belief is, from the epistemic point of view, neither intrinsically valuable nor disvaluable. Likewise, justified or unjustified beliefs about trivial matters are not intrinsically valuable or disvaluable from the epistemic point of view. Consequently, the restricted view cannot account for (F) and (J), because the scopes of both norms are unrestricted. They require us not to have *any* false or unjustified beliefs irrespective of the subject matter. So the perspectival approach must hold *the unrestricted view*: every true (false) belief is intrinsically (dis)valuable from the epistemic point of view. But the difficulty, as we have seen, is that not all beliefs seem intrinsically (dis)valuable even from the epistemic point of view. Hence, Grimm concludes that the perspectival approach faces an unsolvable dilemma: either it has to adopt the unrestricted view in order to make sense of the scope of epistemic normativity, which is implausible, or it can adopt the restricted view, but is unable to make sense of the scope of epistemic normativity.[68]

Another possible response of the value approach is advanced by Ernest Sosa (2007). Sosa's approach avoids the above dilemma for the value approach by holding an unrestrictive view that all true beliefs are valuable *within the epistemic domain*. Sosa maintains that evaluation could be conducted *within a domain*, such as music, sports, food. The evaluation concerns only whether things are valuable within that domain. Indeed, things valuable within a domain may not have any value beyond that domain. As Sosa explains:

> Paradoxically, one can be an adept critic within such a domain even while discerning in it no domain-transcendent value. Thus, someone knowledgeable about guns and their use for hunting, for military ends, and so on, may undergo a conversion that makes the use of guns abhorrent. The good shot is thus drained of any real value that he can discern. Nevertheless, his critical judgment within that domain may outstrip anyone else's, whether gun lover or not. Critical domains

[68] For other criticisms of the value approach, see Selim Berker (2013a; b), Michael DePaul (2001), Thomas Kelly (2003).

> can be viewed as thus insulated, in ways suggested by our example. (Sosa 2007, 73-74)

Epistemic evaluation is also conducted within *the epistemic domain*. It is unclear how to characterise what the epistemic domain is. But epistemic values (namely, values within the epistemic domain) can be classified as either *fundamental* within the epistemic domain or *derivative* insofar as they help to promote or bring out fundamental values. Sosa's approach thus avoids the above dilemma for the value approach. Sosa does not treat epistemic value as a kind of intrinsic value, so he is not committed to the view that we should promote or preserve true beliefs. Instead, we only need to recognise that true belief is a fundamental value within the epistemic domain, just like an accurate or lethal shot is a fundamental value in the military domain. Hence, Sosa's account can make sense of the scope of epistemic evaluation: since all true beliefs are fundamentally valuable within the epistemic domain, justification can also be evaluated as valuable within the epistemic domain, because justified beliefs are more likely to be true. Therefore, Sosa's account offers a plausible explanation as to how all true or justified beliefs are epistemically valuable and all false or unjustified beliefs are epistemically disvaluable.

Howevery, Grimm (2009) convincingly argues that Sosa's account cannot explain why epistemic norms are categorical.[69] The reason why Sosa's account can avoid the dilemma for epistemic teleology is the idea that epistemic values are domain-specific. Sosa's agnosticism about whether epistemic values have any value beyond the epistemic domain renders his approach unable to explain the categoricity of epistemic norms. As Grimm explains:

> What I want to suggest now is that by remaining agnostic about the domain-transcendent value of true belief, Sosa seems to introduce a new problem—seems to, indeed, lose sight of one of the most important aspects of our epistemic appraisals. For notice: when we judge a belief to be unjustified or irrational, we seem to be doing more than just evaluating (in this case, in a negative way) the skill or virtuosity of the believer's performance. In addition, we seem to be in some sense *criticizing*, perhaps even *reproaching*, them for believing in this way.

[69] Pascal Engel also accepts this objection to Sosa: 'One further difficulty that Sosa's performance account of normativity encounters is that a number of normative judgments that we pass on beliefs are of a deontic unrestricted general form, and not an evaluation of the skill or competence of the agent in a particular domain' (Engel 2013b, 622).

> To judge someone's belief to be unjustified or irrational is thus to judge that the person's attitude towards the content of the belief *should* be reconsidered, in some apparently binding sense of 'should.' (Grimm 2009, 253-254)

Basically, Grimm's point is that Sosa's account of epistemic evaluation cannot explain why epistemic norms such as (F) and (J) are categorical. The reason is that since epistemic value, according to Sosa is domain-specific, epistemic values may or may not be valuable beyond the epistemic domain. As Sosa makes clear, a lethal weapon is valuable in the military domain, though people may consider it to have no value at all beyond the military domain. Indeed, many would consider it intrinsically disvaluable, and so may hold that we should not produce or preserve lethal weapons. So, the military value of a lethal weapon cannot generate categorical norms. Similarly, epistemic values, being domain-specific, cannot explain the categoricity of epistemic norms. Grimm concludes that, in order to explain the categoricity of epistemic norms, the value approach must show that epistemic values have *domain-transcendent values*.[70]

I think that Grimm's objection to Sosa is right, but a similar objection can be applied to the perspectival approach. The perspectival approach, as we have seen, faces a dilemma between the unrestricted and restricted views, regarding the value of true belief. This dilemma, however, may not be so serious to the perspectival approach as it appears. Grimm holds that the perspectival approach needs the unrestricted view, but the unrestricted view seems implausible. However, proponents of the perspectival approach can insist that the unrestricted view is correct. They could argue that even trivial true beliefs have some intrinsic value from the epistemic point of view, though their values are easily overridden by the disvalue of contemplating trivial matters.[71] One could certainly argue, as Grimm does, that some true beliefs have no intrinsic value at all even from the epistemic point of view. It is difficult to see which side is right. So Grimm's objection to epistemic teleology is not conclusive.

However, I think that Grimm misses the fact that the perspectival approach faces the same problem as Sosa. Let's assume that the unrestricted view, for the sake of argument, is tenable.

---

[70] Certainly, Sosa could supply an account of the domain-transcendent values of epistemic values. So Grimm's objection is merely that Sosa's value approach is insufficient to explain the categoricity of epistemic norms. However, Sosa (2003) also argues that trivial true beliefs are not intrinsically valuable. Perhaps, Sosa proposes the domain-specific account to avoid that problem.

[71] Paul Horwich (2006), Jonathan Kvanvig (2003), and Michael Lynch (2009a) adopt this argument to defend the unrestricted view. I discuss this argument in greater detail in Chapter 1.

Even so, epistemic values are intrinsically valuable *merely from the epistemic point of view*. The perspectival approach is silent about whether epistemic values are intrinsically valuable when not viewed from the epistemic point of view. So the epistemic norms generated by the perspectival approach would hold that, *from the epistemic point of view*, one should promote epistemic values and prevent epistemic disvalues. (T) and (J) tell us that we should not have false and unjustified beliefs, *period*, rather than that we should not have false and unjustified beliefs, *from the epistemic point of view*. Hence, the perspectival approach does not explain the categoricity of epistemic norms.

Let me consider one response from a proponent of the perspectival approach. She may reply that the epistemic perspective is a perspective which we *must* take. So, even if epistemic norms have force only from the epistemic point of view, since the epistemic perspective is inescapable, we would have to conform to epistemic norms.[72] That is, even if epistemic norms themselves are not categorical, their force is no different from categorical norms. In other words, it is a distinction without difference between the norm that we should not have false or unjustified beliefs, *period*, and the norm that we should not have false or unjustified beliefs, from the epistemic point of view. Therefore, the perspectival approach can grant the point that epistemic norms are not categorical, but holds that it remains true that we, *ceteris paribus*, should not have false or unjustified beliefs.

However, this response misunderstands my criticism. I do not deny that, necessarily, we consider epistemic value as intrinsically valuable from the epistemic point of view, and thus we must admit that, from the epistemic point of view, there are norms requiring us to promote epistemic values and to prevent epistemic disvalues. My point is that the perspectival approach is silent about whether epistemic values are actually intrinsically valuable, *simpliciter*.[73] It is instructive to use a distinction made by Duncan Pritchard between *epistemic value* and *the value of the epistemic*:[74]

---

[72] This response is also available to Sosa's domain-specific approach. My objection is also applicable to Sosa's.

[73] A related but different concern about the perspectival approach is that, if epistemic value is a value merely from the epistemic point of view, how do we weigh it against other kinds of value, such as moral or pragmatic values? Daniel Whiting remarks: 'It is plausible to think that the aim to believe only the truth determines the epistemic perspective. It is unsurprising, then, that appealing to that aim alone should not resolve the puzzle of why that perspective eclipses all others when a subject is faced with the question of what to believe, given that the question appears to be one which can also be asked from the competing practical perspective' (Whiting 2014, 222). Similarly, Linda Zagzebski argues that the perspectival approach cannot satisfy the normativity requirement, because it cannot explain how the

> We cannot (without further argument anyway) infer from the fact that something is valuable (or disvaluable) when assessed along a certain axis of evaluation that it is valuable (or disvaluable) along another axis of evaluation, much less that it is valuable (or disvaluable) *simpliciter*. In particular, for our purposes, it does not follow from the fact that something is lacking in epistemic value that it is lacking in value *simpliciter*. In short, we should recognise the distinction between the value (or disvalue) of the epistemic and epistemic value (or disvalue). (Pritchard 2014)

So, the inescapability of the epistemic perspective implies only that we must consider something as epistemically valuable, but it remains open as to whether epistemic values are intrinsically valuable *simpliciter*. Indeed, the distinction between intrinsic value *simpliciter* and epistemic value (intrinsic value from the epistemic point of view), as we have seen, is essential for the perspectival approach to maintain the unrestrictive view about the epistemic value of true belief: the perspectival approach grants that trivial true beliefs may not be intrinsically valuable *simpliciter*, but insists that they are intrinsically valuable from the epistemic point of view. However, the crucial question is 'Should we promote trivial true beliefs?' The perspectival approach would certainly say that we *should, from the epistemic point of view*. But it is not the case that we *should categorically*, since trivial true beliefs are not intrinsically valuable *simpliciter*. Overall, it is not the case that we should promote trivial true beliefs, even if the epistemic perspective is inescapable. The inescapability of the epistemic perspective, therefore, does not imply that the distinction between categorical norms and norms from the epistemic point of view is a distinction without difference. The perspectival approach remains unable to explain why we should conform to epistemic norms.

To strengthen my point, consider an analogy. Since I am *homo sapiens*, let me also assume that the biological perspective is also inescapable and that my survival and the reproduction of my offspring would be considered intrinsically valuable from the biological perspective.

---

conflicts between values: 'William Alston defines epistemic value as value "from the epistemic point of view". I find this a misleading way to explain our caring about truth, knowledge, understanding, and the epistemic values mentioned in this paper since it encourages the mistake of thinking that epistemic value is a special category of value in competition with, and perhaps incommensurable with, moral and pragmatic value. … If we say that epistemic value is value from an epistemic point of view, moral value is value from a moral point of view, and so on, that tends to stymie discussion because it appears that we have to deal with conflicting points of view in order to resolve the conflict' (Zagzebski 2004, 369-370).

[74] Grimm, in a later paper (Ahlstrom-Vij and Grimm 2013), is also aware of the point that epistemic values may not be valuable *simpliciter* or all-things-considered, though he seems unaware of the point that it undermines the possibility of epistemic evaluation to explain epistemic normativity.

Presumably, the norm, from the biological perspective, is that I should promote my survival and the reproduction of my offspring. Even so, I could still reflect on the value of survival and reproduction and conclude that they are not intrinsically valuable *simpliciter*. Thus, I can conclude that it is not the case that I should promote my survival and the reproduction of my offspring, even though the biological perspective is inescapable.

Therefore, the perspectival approach faces the same problem as Sosa: it cannot explain the categoricity of epistemic norms, because epistemic value is specific to a certain perspective or domain. Epistemic values may not be valuable at all, and hence may offer us no reason to pursue them.

Let me draw some lessons from the discussion in this section. As Grimm argues, to successfully explain the categoricity of epistemic norms, the value approach must propose an account of epistemic value that satisfies the following requirements: first, it must be *unrestricted in scope*; second, it must be *normatively robust enough* in order to generate categorical norms. The problem for the value approach is that if it wants to have a theory of the epistemic value of belief that is unrestricted in scope—a theory that epistemic value is valuable only from the epistemic perspective or within the epistemic domain—the theory would not be normatively robust enough. On the other hand, we see that in the beginning the value approach may propose an account of intrinsic value *simpliciter*. Such an account would be normatively robust enough, but we have seen that it is difficult to maintain that all true beliefs are intrinsically valuable *simpliciter*.[75]

However, it is not hopeless to find something epistemic that is intrinsically valuable *simpliciter* and unrestrictive in scope. I will propose my own value approach in the last section. Before that, I will turn to another prominent evaluative approach to epistemic normativity, namely, the correctness approach. The correctness approach looks promising, but I will argue that it ultimately fails to explain epistemic normativity. Nevertheless, its failure points to the way to the right account.

[75] Could it be the case that all knowledge or understanding is intrinsically valuable *simpliciter*? But if some truths are not intrinsically valuable *simpliciter*, it is difficult to see why knowing or understanding those truths is intrinsically valuable *simpliciter*. See Chapter 1 for more discussion.

5.3 The Correctness Approach

I have argued that the value approach, at least in its fashionable forms, has difficulty in explaining the categoricity of epistemic norms. The value approach only discusses epistemic values or goodness. But evaluative terms are much broader: correctness, rightness, fittingness, virtuousness are arguably evaluative terms.[76] An evaluative approach using those terms, some may argue, could be able to explain epistemic normativity.

Indeed, a popular evaluative approach (Boghossian 2003; Lynch 2009a; McHugh 2014; Shah and Velleman 2005; Wedgwood 2002) is to hold that beliefs are constitutively governed by *the correctness norm*: for example, believing that p is correct if and only if p is true, and incorrect if and only if p is false. Call it *the correctness approach*. How does the correctness approach explain why we should, categorically, conform to epistemic norms such as (F)?[77] Let me use the account developed by Nishi Shah and David Velleman (2005; see also Shah 2003; Evans and Shah 2012) as a representative of how the correct approach would account for (F).

One clue that can be found in the account of Shah and Velleman is that incorrectness implies *impermissibility*, as they say: 'The norm implicit in the concept of belief is that of being correct if and only if true. Holding a false belief is contrary to this norm: it's not permissible, because incorrect' (2005, 512). Perhaps they are offering the following argument for (F):

> (1) If believing falsehood is incorrect, it is impermissible to believe falsehood;
>
> (2) If it is impermissible to believe falsehood, then we should not have false beliefs;
>
> (3) Believing falsehood is incorrect (the correctness norm).
>
> Hence, (F) we should not have false beliefs.

This argument looks plausible. Incorrectness and impermissibility do have a tight connection. If the correctness norm is true, then the correctness approach seems promising with regard to explaining epistemic normativity.

---

[76] Some philosophers may classify those terms differently, which is not my concern here and does not affect my argument.

[77] How about the norm (J)? Many proponents of the correctness approach hold that the justification norm of belief can be derived from the correctness norm. So I focus on the explanation of (F) here.

However, I think that the above argument is wrong, because there are two readings of the correctness norms ('S' stands for a person), as below:

> (STATE) S's belief that p is correct *qua* belief if and only if p is true and incorrect *qua* belief if and only if p is false;
>
> (PERSON) S is correct to believe that p if and only if p is true and incorrect if and only if p is false.

The problem for Shah and Velleman is that while incorrectness and impermissibility are interchangeable in PERSON, they are not interchangeable in STATE. The reason is that what (im)permissibility really governs are agents. Even if we do sometimes say that it is impermissible for an action or a belief to be wrong or that they are impermissible actions or beliefs, what we really mean is that it is impermissible for us to take wrong actions or have incorrect beliefs. So if the PERSON reading of the correct norm is right, then Shah and Velleman can have an explanation of (F). On the other hand, the STATE reading cannot offer what Shah and Velleman want, since correctness and permissibility cannot be used interchangeably in STATE. So whether the argument of Shah and Velleman succeeds depends on whether they can have the PERSON reading of the correctness norm.

Unfortunately for Shah and Velleman, what they actually offer is the STATE reading of the correctness norm, as they claim that it is a conceptual truth about belief that belief is constitutively governed by the correctness norm:

> In our view, being regulated for truth is part of the very concept of belief: to conceive of an attitude as a belief is to conceive of it as a cognition regulated for truth, at least in some sense and to some extent. But we think that the concept of belief must include more than the manner in which the attitude is actually regulated. Also part of the concept is a standard of correctness. Classifying an attitude as a belief entails applying to it the standard of being correct if and only if it is true. (Shah and Velleman 2005, 498)

According to Shah and Velleman, when we conceive our attitude as belief (*qua* belief), we must consider our attitude *qua* belief as governed by the correctness norm. Hence, it is clear that the correctness norm of belief proposed by Shah and Velleman should be read as STATE.[78]

But could the correctness approach argue for PERSON from STATE? Shah and Velleman may respond that we may infer from the truth of STATE to the truth of PERSON when we are deliberating whether we should believe:

> One can deliberate whether to believe that p because one can engage in reasoning that is aimed at issuing or not issuing in one's believing that p in accordance with the norm for believing that p. And one can engage in reasoning with that aim precisely by considering the question *whether* p. Considering *whether* p can amount to reasoning aimed at issuing or not issuing in one's believing that p in accordance with the relevant norm because the relevant norm is this: believing that p is correct if and only if p is true, and hence if and only if p. Furthermore, because the norm is contained in the concept of belief, and doxastic deliberation is framed by the question *whether to believe that p*, anyone engaging in doxastic deliberation knows that this is the relevant norm. (Shah and Velleman 2005, 502-503)

Shah and Velleman argue that, when we deliberate whether we should believe that p, we need to regulate our doxastic deliberation in accordance with the correctness norm. Since truth is the correctness norm of belief, then we should deliberate the question whether p. The procedure of doxastic deliberation, according to Shah and Velleman, may be formulated as the following argument (I focus only on false beliefs, given (F)):

(5) I should not believe that p if the belief that p is incorrect;

(6) Believing that p is incorrect if and only if p is false;

(7) It is false that p.

---

[78] Certainly, proponents of the correctness approach may not consider it a conceptual truth that belief is governed by the correctness norm. Instead, they may regard it as a truth about the nature of belief. This difference does not affect my argument here. Furthermore, proponents of the correctness approach often consider the correctness norm as uncontroversial, even a truism. The PERSON reading of the correctness norm is not uncontroversial. It is disputable why one is impermissible to believe falsehood. Hence, proponents of the correctness approach should adopt the STATE reading and argue from that for the PERSON reading. However, as I will argue shortly, I do not think that such an argument is available.

Hence, (F) I should not believe that p.

(7) is an assumption, and (6) is the correctness norm (though I have not specified upon which reading it is based). The crucial premise is obviously (5). Is (5) true, however? Shah and Velleman claim that doxastic deliberation should be regulated by the correctness norm. But there are two readings of the correctness norm. We have seen that what Shah and Velleman actually provide is STATE. For the present purpose, however, let me assume that Shah and Velleman can have PERSON. They could then prove (5) as follows:

(9) I am incorrect to believe that p if and only if p is false (PERSON);

(10) If I am incorrect to believe that p, I am impermissible to believe that p;

(11) If I am impermissible to believe that p, I should not believe that p.

Hence, (5) if p is false, I should not believe that p.

I hold that (10) and (11) are true, because incorrectness and impermissibility, when applied to agents, can be used interchangeably and can generate categorical norms. Hence, Shah and Velleman can prove (5) if they can have PERSON. The problem is, of course, that they have *not* proved that PERSON is correct.[79] In order to prove (5), Shah and Velleman must prove PERSON, presumably from STATE. Perhaps, they can argue that there is an assumption that bridges the gap between STATE and PERSON:

(STATE) The belief that p is correct *qua* belief if and only if p is true and incorrect *qua* belief if and only if p is false;

(BRIDGE) S is correct to believe that p if and only the belief that p is correct *qua* belief and incorrect if and only if the belief that p is incorrect *qua* belief.

[79] In their paper, Shah and Velleman want to argue that the fact that belief is governed by the correctness norm is the best explanation for the *transparency* of doxastic deliberation: transparency is the feature that the deliberative question *whether I should believe that p* is transparent to *whether p*. As I argue here, the norm that Shah and Velleman should take as governing doxastic deliberation is PERSON. However, transparency is controversial and questioned by some philosophers (McHugh 2011a; 2013; Steglich-Petersen 2006; 2009). So it is unclear to me whether this strategy is successful. To say the least, my discussion highlights the problem that Shah and Velleman fail to distinguish between STATE and PERSON. Even granted that STATE is a conceptual truth about belief, they cannot argue from STATE to PERSON, as I will argue shortly. So all Shah and Velleman have to support PERSON is an argument from best explanation of the transparency of doxastic explanation, which is controversial.

> Hence, (PERSON) S is correct to believe that p if and only p is true and incorrect if and only if p is false.

If BRIDGE is true, then Shah and Velleman can have their cake and eat it too. However, BRIDGE is wrong. To see why, let me turn to other kinds of mental state. Shah (2008; see also Evans and Shah 2012) argues that intention and desire are also governed by their own correctness norms. Judith Thomson (2008, ch.7) holds that many kinds of mental state—such as trusting, disliking, admiring, despising—are also subject to some correctness norms. Whether their views are correct is irrelevant to my concern here. My concern is that if those mental states are also governed by some correctness norms, there is no BRIDGE* that connects STATE* and PERSON* (asterisks are added in order to signify that they are applied to states in general, rather than to belief).

Consider despising. Suppose despising is governed by STATE*: despising x is correct *qua despising* if and only if x is despicable, and incorrect *qua* despising if and only if x is not despicable. Presumably, Adolf Eichmann was despicable, so that despising him is correct *qua* despising. If BRIDGE* is correct, then one can hold PERSON*: S is correct to despise x if and only if x is despicable, and incorrect to despise x if and only if x is not despicable. It follows that we are correct to despise Adolf Eichmann.

However, I think that this is wrong. Many people may have an ethical or religious belief according to which despising people is intrinsically bad. Jesus and Mozi taught people that they should forgive and love their enemies. Buddha preached that people should get rid of negative emotions. Their followers would grant that despising Eichmann is correct *qua despising*, but reply that we are not correct to despise him, because despising people is intrinsically bad. The fact that despising Eichmann is correct *qua despising* does not imply that we are correct to despise him. BRIDGE* should be rejected.

One may respond that my conclusion from the example is wrong, because it could be the case that, *other things being equal*, we are still correct to despise Eichmann. However, we also need to conform to social, ethical, or religious norms or values. Very often, the ethical or religious concerns outweigh the concern for correctness of mental state. The correctness approach can argue that it remains true that, other things being equal, we are still correct to despise Eichmann.

However, I think that this response misinterprets the case. For Christians, Buddhists, or Mohists, their view is not just that the ethical or religious concern should, perhaps necessarily, outweigh the concern for correctness of mental state. Rather, their view is that one should never despise people. As a theoretical question whether despising Eichmann is correct *qua* despising, they may agree that it is correct *qua* despising, but would insist that it is in fact irrelevant to whether one is correct to despise Eichmann. Despising people is intrinsically bad (or violating God's commands), no matter whether it is correct or incorrect *qua* despising. So they reject BRIDGE* that one can infer from the claim that one's despising is correct *qua* despising to the claim that one is correct to despise.[80]

To be clear, my argument does not rest on whether those ethical or religious teachings are correct. The example is used to illustrate the point that the correctness of a mental state does not imply that it is correct for one to be in that state. The moral from the example is that being in a correct mental state could be intrinsically bad. If something is intrinsically bad, as I discussed above, it means that, *ceteris paribus*, one should stop it happening. Hence, if despising people is intrinsically bad, *ceteris paribus*, one is incorrect to despise people, even those despicable ones. Similarly, being in an incorrect mental state could be intrinsically good or simply value-neutral. If it is value-neutral, one is permissible to be in that state. If it is intrinsically good, *ceteris paribus*, one should be in that state. In other words, BRIDGE* is wrong.[81]

Could proponents of the correctness approach respond that there is something peculiar about belief that makes a principle like BRIDGE possible? If the above reasoning is on the right track, the prospect of this response is dim. Correct (true) beliefs may not be intrinsically valuable at all or even bad, and incorrect (false) beliefs may be intrinsically good or value-neutral. One can certainly try to argue that all true beliefs are intrinsically valuable and all false beliefs are intrinsically bad. However, that would mean giving up the correctness approach and bringing us back to the value approach, which, as I have argued, is problematic.

---

[80] One may respond that, for them, no human is despicable, so they should never despise anyone. However, I think that this response still misinterprets their views. They do not deny that some humans are bad. What they deny is that we can treat them badly.

[81] My point here is similar to, though in different contexts, that which Clayton Littlejohn tries to make (he attributes this point to Judith Thomson (2008)): 'Even if [the knowledge norm] is the constitutive rule of assertion, it doesn't seem to follow that we're under some obligation not to assert what we don't know. … it looks as if there might be two kinds of epistemic norm, those that specify constitutive rules and those that determine our epistemic obligations' (Littlejohn 2014, 2).

To recap, I have argued that proponents of the correctness approach confuse between STATE and PERSON. The correctness approach begins with STATE, which is quite uncontroversial. In order to explain epistemic normativity, however, it needs PERSON. Proponents of the correctness approach seem to assume BRIDGE to connect STATE and PERSON. It seems natural to infer that since one's mental state is correct *qua* its own kind, one is correct to be in that state. However, I have argued that we should be sceptical about BRIDGE. Without an argument for PERSON, the correctness approach has not successfully explained epistemic normativity.

5.4 The Personal Value Approach

In this section, I will propose my own account of epistemic normativity. Before that, let me recapitulate what I have argued so far. I have argued that two prominent evaluative approaches to epistemic normativity are both unsuccessful: the value approach is unable to do so, because it cannot explain how the epistemic value of truth is intrinsically valuable *simpliciter*; meanwhile the correctness approach is also unable to do so, because it confuses the correctness norm that applies to mental states and the correctness norm that applies to persons. One reason why the correctness norm of mental states does not entail the correctness norm for a person is that being in a correct state may not have intrinsic value and being in an incorrect state may not have intrinsic disvalue. In sum, both approaches share the same structural problem: they both focus on the evaluation of *belief*, which, as Grimm points out, must be unrestricted in scope. It seems not to be the case that all true beliefs are intrinsically valuable *simpliciter* and all false beliefs are intrinsically disvaluable *simpliciter*. So, proponents of the evaluative approach may turn to the perspectival or domain-specific approach, or to the correctness approach: both approaches may plausibly hold that *all* true beliefs are epistemically valuable or correct *qua* belief. However, the problem is that epistemic value, so understood, or correctness *qua* belief, is simply not normatively robust enough to explain the categoricity of epistemic norms such as (T) and (J). The real dilemma for the evaluative approach, therefore, is that it cannot find a kind of evaluative feature of belief that is unrestricted in scope and normatively robust enough.[82]

[82] In Chapter 1, I propose an account of epistemic value in terms of goodness of epistemic kind. Accordingly, the epistemic value of true belief is the value of being good *qua* belief. My account of the epistemic value of true belief is unrestricted in scope. However, it is not normatively robust enough to generate norms such as (F) or (J). The reason is that the goodness of a kind does not offer reasons to promote it. It only requires us to appreciate it as being good of its own kind. Hence, it is not normatively robust enough. So, an evaluative approach based on my account of epistemic value of true belief would

Epistemic Evaluation and Epistemic Normativity

In the last section, I distinguished between state-level and person-level correctness. As I have argued, if the correctness approach can prove personal-level correctness, it can explain the categoricity of epistemic norms. So I would like to suggest that epistemic normativity is grounded in the epistemic evaluation of *person*. Call it *the personal value approach*. The personal value approach bypasses the problem about the value of true belief. More importantly, I will suggest that the personal value approach can avoid the dilemma facing the value approach and can explain the categoricity of epistemic norms.

Epistemic value is the central notion in a value approach, so let me begin with the notion of epistemic value. I have argued in Chapter 1 that epistemic value should be understood as *goodness of epistemic kinds*. Epistemic kind is a species of evaluative kind that sets its own evaluative standards, and the evaluative standards of an epistemic kind are truth-directed. For example, thermometer is an epistemic kind, because it sets its own evaluative standards, and its standards are related to truth: a thermometer is good *qua* thermometer if its temperature reading is reliable and accurate.[83]

Person is also an evaluative kind. Some of its evaluative standards are not truth-directed, but some are, which makes person an epistemic kind. Let me explain the idea of person as an evaluative kind first, and as an epistemic kind later.

A theory that elaborates the evaluative standards of being good *qua* person is a theory of the well-being or flourishing of person. The idea is simple: what makes us better people constitutes our well-being or flourishing *qua* person. That is, what is *good for us qua person* constitutes our well-being as people. Hence, the notion of 'X is good for S *qua* K' is used in the *welfarist sense* about the contribution of X to the well-being of S *qua* K. In the literature on the welfarist notion of goodness, the phrase '*qua* K' is often dropped, and philosophers usually focus only on the notion of 'X is good for S' (Rosati 2006; 2009; Zimmerman 2009; Fletcher 2012). But this difference is merely superficial. For while 'S' is always a singular term, other philosophers may replace it with a general term. To express the idea that happiness contributes to the well-being of person, for example, those philosophers may phrase it as 'being happy is good for people', but I prefer 'being happy is good for S (an individual) *qua* person'. The difference might look insignificant. However, those philosophers may sometimes

---

not work. However, my goal of this section is to show how my account of epistemic value can nonetheless provide the basis of epistemic normativity if we focus on the evaluation of person.

[83] The example refers to David Armstrong's classic thermometer model of knowledge (1973).

consider questions like 'what is good for me' without specifying what kind I belong to. Usually when the question 'what is good for me (or an individual person)' is asked, in most contexts, the question is about what is good for me *qua* human or person. However, it is often useful to specify which K is at issue. For example, going to Tahiti was good for Gauguin *qua* painter, but not good for him *qua* husband and father. The following question may then be 'Was being a good painter or being a good husband and father good for Gauguin *qua* person?' Without specifying which K is at issue, it is difficult to respond to the question 'Was going to Tahiti good for Gauguin?'

Furthermore, we can distinguish two senses in which X is good for S *qua* K: first, X is *finally* good for S *qua* K if and only if X is an evaluative standard of K; second, X is *instrumentally* good for S *qua* K if and only if X is not an evaluative standard of K, but helps S satisfy the goodness conditions of K. For example, an evaluative standard of swimmer is having strong muscles. Having strong muscles is thus finally good for me *qua* swimmer, while doing weight training is instrumentally good for me *qua* swimmer, because weight training can help me build stronger muscles, but is not itself an evaluative standard of swimmer.

Let's turn to person as an epistemic kind. Some evaluative standards of person are truth-directed, because a person can do cognitive performances and have many cognitive apparatuses. I use 'cognitive apparatus' very broadly to refer to cognitive abilities (such as perceptual or intellectual abilities), personal traits or habits (such as open-mindedness, carefulness, bias), skills (such as computer or mathematical skills or skills of using laboratory equipment), etc. The kinds of cognitive apparatus and cognitive performance are also epistemic kinds, if their evaluative standards are truth-directed.

My account of epistemic evaluation is a modification of Sosa's virtue-theoretic account (Sosa 2007; 2011; see also Greco 2010). Sosa regards epistemic evaluation as a kind of performance evaluation. A performance can be evaluated based on whether it is successful and competent.[84] For example, when an archer takes a shot, his shot is successful if and only if the shot hits the target, and his performance is competent if he takes the shot skilfully. Similarly, according to Sosa, a belief is a cognitive performance: successful if true, and competent if adroit or skilful.

[84] In Sosa's account, a performance can also be evaluated as *apt* (successful because competent). However, as I have argued in Chapter 4, aptness does not make a performance more valuable. Therefore, I will not discuss aptness in this chapter.

Unlike Sosa, who takes belief as the focal point of epistemic evaluation, the focus of my personal value approach is the performer, i.e., the person.[85] A person is epistemically good *qua* person, if and only if (i) she is epistemically competent *qua* person, or (ii) she delivers many successful cognitive performances and has more successful cognitive performances than unsuccessful ones, that is, she has many true beliefs and the ratio of her true beliefs over false belief is high.

(i) should be manifestly true, though some clarifications may be needed. Very roughly, one is epistemically competent, if and only if one's cognitive apparatus is competent. Nevertheless, one can be competent in an apparatus A, but incompetent in an apparatus B. For example, Beethoven, after becoming deaf, was incompetent in hearing, but competent in sight-reading music.

(ii) is more controversial. Some may object that a person could luckily obtain many true beliefs even though she is utterly incompetent. They may insist that only true beliefs obtained by her own competence can make her epistemically good *qua* person. However, it seems to me that one is an epistemically good person when one's true beliefs reach critical mass and the ratio of one's true beliefs over false beliefs is high. Moreover, storing beliefs is arguably a cognitive capacity, which can be evaluated by the number of beliefs of storage and the ratio of true beliefs over false beliefs. If so, (ii) can be reduced into (i). So I will grant that person is an epistemic kind.

The fact that person is an epistemic kind implies that being epistemically good *qua* person is finally good for us *qua* person. It means that becoming an epistemically better person makes us become better persons. The epistemic goodness of person contributes to our well-being or flourishing *qua* person. The reason is simple: since a theory of goodness *qua* person is a theory of the well-being or flourishing *qua* person and the epistemic goodness of person is included in the theory of goodness *qua* person, the epistemic goodness of person makes us more flourishing *qua* person or become better persons. I consider it self-evident, but some clarification might be needed.

---

[85] In Chapters 2-4, I argue for the performance-based account of epistemic evaluation. The performance-based account and the personal value approach are in the same spirit: both require that epistemologists shift their attention from belief to person (the performer). The performance-based conception holds that part of epistemic evaluation is about the epistemic competence of the person. The personal value approach holds that what grounds epistemic normativity is the fact that being epistemically competent is finally good for us *qua* person.

To begin with, given my account of the epistemic goodness of person, it would follow that having excellent vision—a kind of epistemic competence—contributes to our goodness *qua* person. Some might think that the idea sounds odd. Do I mean that one who has excellent vision is a good person? Or even that a blind person is a bad person? True, I use 'being good *qua* person' and 'being a good person' interchangeably. However, 'being a good person' often means 'being a *morally* good person'. And I certainly do not mean that having excellent vision makes us morally good persons, or that blind people are morally bad people. When I use 'being good *qua* person' or 'being a good person', I mean 'being in a good condition of being a person'. Hence, a theory of goodness *qua* person is a theory of the well-being *qua* person, because well-being *qua* person simply means a good or excellent condition of being a person. Accordingly, there is nothing odd to say that blindness harms our well-being *qua* person, or that having excellent vision makes us better *qua* person.

Moreover, saying that having excellent vision is good for us *qua* person does not mean that people who have excellent vision are good persons, because they might not be good in other respects as persons. However, it remains true that having excellent vision improves the condition of being a good person. Consider an analogy. A dysfunctional car that has a good engine is certainly not a good car, but it remains true that having a good engine is good for it *qua* car. Besides, having excellent vision and being intellectually virtuous are both good for us *qua* person. However, I do not mean that they are *equally* good for us *qua* person, but only that each can more or less contribute to our well-being *qua* person. Perhaps being intellectually virtuous contributes to our personal well-being significantly more than having excellent vision. Or perhaps their contributions cannot be commensurate. For the present purpose, I do not need to decide this. All I need is the point that being epistemically good *qua* person contributes to our well-being or flourishing *qua* person.

Let us return to the personal value approach to epistemic normativity. The categoricity of epistemic norms, I will argue, should be grounded in the epistemic value of person. The reason why the epistemic evaluation of belief fails to explain epistemic normativity, as we have seen, is that a good or correct belief may not be intrinsically valuable at all. To remedy this problem, my personal value approach needs to explain why being an epistemically good person is intrinsically valuable.

Before proposing my own answer, I want to first discuss Jason Baehr's account (2011, ch.6), which offers an account of the intrinsic value of an epistemically good person. The value of

being an epistemically good person, in Baehr's term, is *personal intellectual worth*: it is 'a reasonably familiar and intuitive way of being good or excellent *qua* person that is distinctively intellectual' (2011, 94). Like me, Baehr also holds that a person becomes an epistemically better person if she possesses intellectually good traits, i.e., intellectual virtues (such as open-mindedness, honesty, fair-mindedness).[86] Baehr goes on to explain why intellectual personal worth is valuable through an account of the basis of personal worth:

> (BPW) A subject S is good or better *qua* person to the extent that S is positively oriented toward or 'love' what is good and is negatively oriented toward or 'hates' what is bad. (Baehr 2011, 97)

Baehr's account is developed from the accounts of Robert Adams (1999) and Thomas Hurka (2001). Both hold that, if x is intrinsically good, valuing x for its own sake is also intrinsically good.[87] Therefore, Baehr holds that the value of personal worth is *derivative* (but still *intrinsic*) from the intrinsic goods through the person's intrinsic love of them.

I agree with Baehr that there could be derivative intrinsic values. For Baehr's account to succeed, however, he must show that there are *non-derivative intellectual* intrinsic goods, from which our intrinsic love of them derives its intrinsic value. What are the intellectual intrinsic goods? Baehr answers: 'It will suffice to think of the relevant goods or values in terms of what Linda Zagzebski has called "cognitive contact with reality," a notion that is intended to encompass a range of familiar cognitive states like true belief, knowledge, and understanding' (Baehr 2011, 101). However, this answer immediately raises a difficulty we are familiar with, namely, that some true beliefs or knowledge are not intrinsically valuable. Indeed, Baehr, in his later article (2012), argues that some knowledge and true beliefs are intrinsically worthless or bad, such as truths or knowledge about the latest political sex scandals or the private lives of celebrities.[88] Baehr must concede that intrinsically loving the truth about gossip and scandal is

---

[86] One difference between my account and Baehr's is that Baehr does not think that cognitive faculties or abilities contribute to personal intellectual worth. He considers them as impersonal or non-personal, because the evaluations of cognitive faculties and abilities are 'not evaluations of person *qua* persons. Rather, as the language just employed suggests, it may be more accurate to characterise them merely as evaluations of certain capacities *of* persons (eyesight, memory, etc.)' (Baehr 2011, 93). However, I don't think that this reasoning is correct. After all, one can also evaluate intellectual traits merely as traits of persons. Certainly, one way to evaluate a person is to evaluate her capacities, just like one way to evaluate a cook is to evaluate his culinary skills.

[87] Kurt Sylvan (2012) also develops this line of thought in his account of epistemic values.

[88] How about understanding? It is unclear to see whether understanding is intrinsically valuable and even if it is, how it helps to account for epistemic normativity, because it remains disputed what

intrinsically bad or worthless. Baehr's account, therefore, gets into the same difficulty of the value approach as that which Grimm points out: it cannot offer an account of epistemic intrinsic value that is unrestricted in scope. Accordingly, it cannot explain epistemic norms such as (F) and (J), which are also unrestricted in scope.

Baehr is wrong, I think, in holding that the intrinsic value of personal worth is derivative from other intrinsic goods. Being good *qua* person is a non-derivative intrinsic good. I do not deny that loving intrinsic goods is—derivatively or non-derivatively—intrinsically good, but the intrinsic value of goodness *qua* person is not based on that. Indeed, it is difficult to understand why there is a need to find the basis for the intrinsic value of goodness *qua* person. Being a good person, obviously, is one of the most important and valuable things. It is not obvious that truth and knowledge are more valuable than the goodness *qua* person, or why the value of goodness *qua* person needs to be based on truth and knowledge. On the contrary, it seems closer to the truth that the value of truth and knowledge is derivative from the value of being good *qua* person (in a way I will explain later). So, the difference between Baehr and me, in this regard, is that our explanations of the value of personal goodness run in opposite directions: Baehr claims that the value of personal goodness is derived from the intrinsic value of truth or knowledge, whereas I take the value of personal goodness as fundamental and claim that the intrinsic value of truth and knowledge is derivative. Furthermore, my account can avoid the scope problem of Baehr's account: there is no one's personal goodness that is not intrinsically valuable. Since person is an epistemic kind, all truth and knowledge would be intrinsically valuable for being finally good for us *qua* person. Let me explicate this in the following.

First, is the goodness *qua* person intrinsically valuable? I mentioned above that philosophers often use INTENTIONAL as a test for intrinsic value: if something is evaluated by itself as intrinsically valuable, we should, *ceteris paribus,* promote it, preserve it, or bring it about. I think that there is no doubt that the goodness of a person *qua* person passes the test. It is an ideal that we become good persons, and it is intrinsically bad to hinder one's becoming a better person. Moreover, since INTENTIONAL is an isolation test that considers only the object under evaluation, we should promote personal goodness for its own sake rather than for the

---

understanding is and how it differs from knowledge. Some (Elgin 2009; Riggs 2009) argue that understanding is not factive. If that is the case, why is there the epistemic requirement that we should not have false beliefs? Moreover, if knowledge of the latest sex scandals is intrinsically bad, it seems that understanding of them would be similarly bad.

sakes of others. So the intrinsic value of personal goodness is not derivative. Admittedly, claims about intrinsic value often remain debatable. But I think it is clear that if there is anything intrinsically valuable, personal goodness is a promising candidate. I doubt that anyone would deny that, other things being equal, we should help people, including ourselves, become better people. Hence, I conclude that all personal goodness is non-derivatively intrinsically valuable.

Since personal goodness is intrinsically valuable, epistemic personal goodness is also intrinsically valuable, because epistemic personal goodness constitutes personal goodness. It follows from INTENTIONAL that there is a categorical norm to which we should conform:

(P) We should, *ceteris paribus*, be epistemically good people.

(P), I will explain, is the basis of the normativity of the norms like (F) and (J). The reason is thus: to be an epistemically good person, we need to satisfy the criteria of the epistemic personal goodness I discussed above: (i) epistemic competence *qua* person, and (ii) the quantity and the ratio of true beliefs over false beliefs. That means that truth and justification are also intrinsically valuable, but their intrinsic values are derivative from their contribution to epistemic personal goodness. Having true beliefs and avoiding false beliefs, obviously, help one become an epistemically better person by directly satisfying (ii). As to justification, justification is often explained in terms of truth-conduciveness. Having one's beliefs justified is instrumental to increasing the ratio of true beliefs over false beliefs. But this only shows that justification is merely instrumentally valuable. Nevertheless, if one believes unjustifiably, one is epistemically incompetent for having beliefs by carelessly reasoning from limited evidence, uncritically accepting information from unreliable sources, etc. So justification derives its intrinsic value by satisfying (i). That is, epistemic justification manifests one's epistemic competence and thus makes one an epistemically good person. In sum, true beliefs and justified beliefs are intrinsically valuable by increasing epistemic personal goodness, and false beliefs and unjustified beliefs are intrinsically disvaluable by reducing epistemic personal goodness. Since true beliefs and justified beliefs are intrinsically valuable and false beliefs and unjustified beliefs are intrinsically disvaluable, it follows from INTENTIONAL that we should, *ceteris paribus*, promote true beliefs and justified beliefs, and that we should, *ceteris paribus*, prevent false beliefs and justified beliefs. That is the reason why (F) and (J) are categorical.

An upshot of my account is that it may allow that we are permissible to have a false or unjustified belief if having that belief, *all things considered*, helps one become better *qua*

person. For example, Sarah is diagnosed that she has at most three months to live. Her daughter is pregnant, and is due in six months. Sarah would love to see her grandchild and would feel bereft if she thinks that she would not be able to do so. Believing that she could live to see her grandchild would give her great comfort and strength for the rest of her life. Those factors considered, she should believe that she will see her grandchild. In doing so, she may become a less epistemically good person, but living happily and with hope is also good for her *qua* person. Epistemic normativity is thus categorical, but it is *pro tanto*.

So, my personal value approach can explain epistemic normativity. By focusing on the value of epistemic personal goodness, my approach is able to offer an account of epistemic value that is normatively robust enough and unrestricted in scope. Hence, my approach solves the dilemma facing traditional evaluative approaches which focus on belief. Now let me consider one response. One may argue that my explanation still relies on the claim that all true beliefs are intrinsically valuable. However, as I have repeatedly argued, other value approaches fail, because they cannot maintain that all true beliefs are intrinsically valuable. Am I inconsistent? No. The reason is that the intrinsic value of true belief, according to my approach, is derivative from the intrinsic value of epistemic personal goodness. However, other approaches simply hold that they are non-derivatively intrinsically valuable. When true beliefs are evaluated by themselves, it is not plausible that they have non-derivative intrinsic value. On the other hand, if we consider the contribution of true beliefs to epistemic personal goodness, even trivial true beliefs can be intrinsically valuable.[89] So, my objection to other approaches is not that some true beliefs are not intrinsically valuable, but that those approaches fail to make sense of why those seemingly worthless beliefs can have intrinsic value.

In conclusion, I have offered an explanation of epistemic normativity. Epistemic normativity is grounded in epistemic personal goodness, which is, in turn, grounded in personal goodness. Personal goodness is non-derivatively intrinsically valuable, so we should promote personal goodness. One interesting consequence of my approach is that there is no strict distinction between epistemic normativity and ethical normativity. The traditional evaluative approaches usually hold that epistemic normativity is distinct from ethical normativity, because it is grounded in values from the epistemic point of view, or in the nature or concept of belief (or in theoretical or doxastic deliberation). However, according to my approach, epistemic

[89] This shows that the isolation tests like INTENTIONAL are tests only for non-derivative intrinsic values. So trivial true beliefs, considered independently, cannot pass the tests. However, it remains possible that they are intrinsically valuable for their contribution to personal goodness.

normativity is grounded in personal goodness. The reason why we should conform to epistemic norms is that doing so makes us become better persons. Epistemic personal goodness is a constituent of personal goodness. To that extent, epistemic evaluation is a subset of ethical evaluation, so is epistemic normativity a branch of ethical normativity.

Conclusion

Let me summarise the thesis. The first chapter was on epistemic value. I argued that the orthodox account of epistemic value—according to which epistemic value is a kind of intrinsic value from the epistemic point of view—should be rejected. My reasons were that (i) the orthodox account is the cause of the triviality objection and (ii) it encourages epistemic teleology, which suffers from the consequentialist problem. I proposed a new account of epistemic value according to which epistemic value is a kind of final value (the goodness of an epistemic kind). I then showed that my account can solve the triviality objection and the consequentialist problem.

In the next three chapters, I attempted to combine my account of epistemic value with the virtue-theoretic account of epistemic normativity in order to offer a satisfactory account of the value of knowledge. The virtue-theoretic account considers epistemic evaluation as a species of performance evaluation. It immediately gives rise to the problem of how a mental state, such as belief or knowledge, can be evaluated as a performance. I showed, in Chapter 2, how performance evaluation is applicable to states. The key idea is that some states are the results of performances and the results of performances can be considered as performances in a derivative sense. That is, when a result of a performance is evaluated in terms of the standards of performance evaluation, what is really under evaluation is the performance that achieves the result: whether the performance *successfully* achieves the result it aims at, whether the performance is *competent*, and whether the performance is apt (successful because competent). This point is crucial to my account of the value of knowledge.

In Chapter 3, I argued that the swamping problem is caused by the assumption that the value of knowledge consist entirely of the values of belief. I call it 'the belief-based conception of epistemic evaluation'. The belief-based conception considers knowledge as a kind of belief, so the swamping problem is read as a problem about how knowledge is a better belief *qua* belief than mere true belief. To solve the swamping problem, according to the belief-based conception, one must find a property of belief that can confer addition value to true belief. I

used Linda Zagzebski's coffee analogy to highlight this point. Just like there is no property of coffee that can make a tasty coffee better *qua* coffee, there is no property of belief that can make a true belief better *qua* belief. Based on my account of the virtue-theoretic view in Chapter 2, I argued that we should adopt the performance-based conception of epistemic evaluation, which holds that the value of knowledge can consist in the value of cognitive performance. The value of cognitive performance is finally valuable, so the swamping problem is solved.

In Chapter 4, I criticised Greco and Sosa for their view that being apt makes a performance better *qua* performance. I argued that the value of performance *qua* performance consists entirely of the value of success and the value of competence. Hence, given the virtue-theoretic account of knowledge, knowledge is no more valuable than gettierised belief. I suggested that virtue epistemologists should turn their attention to epistemic competence.

In Chapter 5, I turned to epistemic normativity. I argued that epistemic normativity cannot be grounded in the evaluation of belief. The reason is that the value of being epistemically good or correct *qua* belief is not normatively robust enough to explain the categoricity of epistemic norms. Instead, epistemic normativity is grounded in the evaluation of a person, because the goodness of a person *qua* person is intrinsically valuable, which is normatively robust enough to generate categorical norms.

There are some themes in this thesis that I would like to explore further in the future.

The first one is metaepistemological. If my account of epistemic value and epistemic normativity is correct, it follows that there is nothing special or distinctive about them. Evaluating the value of belief *qua* belief is just like evaluating the value of knife *qua* knife. Indeed, we can just drop the term 'epistemic value'. Epistemic normativity, as I have argued, is a subset of ethical normativity. The reason why we should not have false or unjustified beliefs is exactly the same as the reason why we should not take wrong actions: that is, to prevent us becoming bad *qua* person. In my thesis, however, I only touched upon some approaches to epistemic normativity. I would like to explore other approaches to examine the nature of epistemic normativity.

The second one is ethical. My account of epistemic normativity relies on the idea of personal goodness. In the thesis, I assumed that person is an evaluative kind, and I did not fully spell out the evaluative standards of person *qua* person. A theory of personal goodness is beyond the

scope of this thesis. Since it is the pivot of my picture of normativity, this topic will be one of my central concerns.

The third one is value theory. I argued that epistemic value is better to be understood as goodness of epistemic kinds. Therefore, I regard epistemic value as relative goodness (attributive goodness), instead of absolute (predicative, intrinsic) goodness. In Chapter 5, however, I argued that the ground of epistemic normativity is the intrinsic value of epistemic personal goodness. In other words, I was suggesting that some kinds of attributive goodness are absolutely or intrinsically good. In value theory, philosophers have long debated about which kind of goodness is more central to our ethical reasoning.[90] Absolutists argue that what is good for a person should be explained in terms of intrinsic values and their relations to the person (usually in his possession), whereas relativists argue that the talk of intrinsic goodness should be ditched from ethical reasoning. I think that it is a mistake to think that absolute goodness and relative goodness are competing at the same level. Instead, I think that relative goodness is a lower-level one concerning the conditions for the well-being of an object *qua* its own kind, while absolute goodness is a higher-level one concerning the ethical status of relative goodness. For example, the goodness of knife *qua* knife and the goodness of person *qua* person are both good relative to the standards of their own kinds. However, the goodness of person *qua* person is absolutely or intrinsically good, but the goodness of knife *qua* knife is not. The idea of absolute or intrinsic goodness signifies the fact that the goodness of person *qua* person enjoys a full ethical status, but the goodness of knife *qua* knife does not. I think that this is a promising way to solve the debate, which I plan to develop in future.

The final one is virtue theory. I argued that virtue epistemologists should turn their attention to epistemic competence and that personal goodness is the ground of epistemic normativity. Virtue is certainly a type of competence and contributes to our personal goodness. So I would like to study virtue theory and individual virtues.

[90] For an recent example of this debate, see Christian Piller (2014).

Bibliography

ADAMS, R. M. 1999. *Finite and Infinite Goods: A Framework of Ethics,* Oxford, Oxford University Press.

AHLSTROM-VIJ, K. & DUNN, J. 2014. A Defence of Epistemic Consequentialism. *The Philosophical Quarterly,* 64**,** 541-551.

AHLSTROM-VIJ, K. & GRIMM, S. R. 2013. Getting it Right. *Philosophical Studies,* 166**,** 329-347.

ALEXANDER, J. & WEINBERG, J. M. 2007. Analytic Epistemology and Experimental Philosophy. *Philosophy Compass,* 2**,** 56-80.

ALSTON, W. 1989. *Epistemic Justification: Essays in the Theory of Knowledge,* Ithaca, Cornell University Press.

ALSTON, W. P. 1988. The Deontological Conception of Epistemic Justification. *Philosophical Perspectives,* 2**,** 257-299.

ALSTON, W. P. 1993. Epistemic Desiderata. *Philosophy and Phenomenological Research,* 53**,** 527-551.

ALSTON, W. P. 2005. *Beyond Justification: Dimensions of Epistemic Evaluation,* Ithaca, Cornell University Press.

ARMSTRONG, D. M. 1973. *Belief, Truth, and Knowledge,* Cambridge, Cambridge University Press.

BAEHR, J. 2009. Is There a Value Problem? *In:* HADDOCK, A., MILLAR, A. & PRITCHARD, D. (eds.) *Epistemic Value.* Oxford: Oxford University Press.

BAEHR, J. 2011. *The Inquiring Mind: On the Intellectual Virtues and Virtue Epistemology,* Oxford, Oxford University Press.

BAEHR, J. 2012. Credit Theories and the Value of Knowledge. *The Philosophical Quarterly,* 62**,** 1-22.

BATTALY, H. 2008. Virtue Epistemology. *Philosophy Compass,* 3**,** 639-663.

BERKER, S. 2013a. Epistemic Teleology and the Separateness of Propositions. *Philosophical Review,* 122**,** 337-393.

BERKER, S. 2013b. The Rejection of Epistemic Consequentialism. *Philosophical Issues,* 23**,** 363-387.

BERKER, S. 2015. Reply to Goldman: Cutting Up the One to Save the Five in Epistemology. *Episteme,* 12**,** 145-153.

BODANSZKY, E. & CONEE, E. 1981. Isolating Intrinsic Value. *Analysis,* 41**,** 51-53.

BOGHOSSIAN, P. A. 2003. The Normativity of Content. *Philosophical Issues,* 13**,** 31-45.

BONJOUR, L. 1985. *The Structure of Empirical Knowledge,* Cambridge, MA., Harvard University Press.

Bibliography

BRADLEY, B. 2006. Two Concepts of Intrinsic Value. *Ethical Theory and Moral Practice,* 9**,** 111 - 130.

BRENTANO, F. 1969. *The Origin of Our Knowledge of Right and Wrong,* Abingdon, Routledge.

BROAD, C. D. 1930. *Five Types of Ethical Theory,* London, Routledge and Kegan Paul.

BROGAARD, B. 2006. Can Virtue Reliabilism Explain the Value of Knowledge? *Canadian Journal of Philosophy,* 36**,** 335-354.

CARTER, J. A. forthcoming. Robust Virtue Epistemology As Anti-Luck Epistemology: A New Solution. *Pacific Philosophical Quarterly*.

CARTER, J. A. & JARVIS, B. 2012. Against swamping. *Analysis,* 72**,** 690-699.

CARTER, J. A., JARVIS, B. & RUBIN, K. 2013. Knowledge: Value on the Cheap. *Australasian Journal of Philosophy,* 91**,** 249-263.

CHISHOLM, R. 1989. *Theory of Knowledge,* Englewood Cliffs, Prentice Hall.

CHISHOLM, R. M. 1978. Intrinsic Value. *In:* GOLDMAN, A. I. & KIM, I. (eds.) *Values and Morals.* Boston: D. Reidel.

CHISHOLM, R. M. 1981. Defining Intrinsic Value. *Analysis,* 41**,** 99 - 100.

CHRISMAN, M. 2012. The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions. *Journal of Philosophy,* 109**,** 588-612.

CHUARD, P. & SOUTHWOOD, N. 2009. Epistemic norms without voluntary control. *Noûs,* 43**,** 599-632.

COLAÇO, D., BUCKWALTER, W., STICH, S. & MACHERY, E. 2014. Epistemic Intuitions in Fake-Barn Thought Experiments. *Episteme,* 11**,** 199-212.

CRAIG, E. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis,* Oxford, Oxford University Press.

DAVID, M. 2001. Truth as the Epistemic Goal. *In:* STEUP, M. (ed.) *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue.* Oxford: Oxford University Press.

DAVID, M. 2014. Truth as the Primary Epistemic Goal. *In:* STEUP, M., TURRI, J. & SOSA, E. (eds.) *Contemporary Debates in Epistemology.* 2 ed. Oxford: Wiley-Blackwell.

DEPAUL, M. R. 2001. Value Monism in Epistemology. *In:* STEUP, M. (ed.) *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue.* Oxford: Oxford University Press.

DUTANT, J. 2013. In Defence of Swamping. *Thought: A Journal of Philosophy,* 2**,** 357-366.

ELGIN, C. 2006. From Knowledge to Understanding. *In:* HETHERINGTON, S. C. (ed.) *Epistemology Futures.* Oxford: Oxford University Press.

ELGIN, C. 2009. Is Understanding Factive? *In:* HADDOCK, A., MILLAR, A. & PRITCHARD, D. (eds.) *Epistemic Value.* Oxford: Oxford University Press.

ENGEL, P. 2013a. Is Epistemic Agency Possible? *Philosophical Issues,* 23**,** 158-178.

ENGEL, P. 2013b. Sosa on the Normativity of Belief. *Philosophical Studies,* 166**,** 617-624.

EVANS, M. & SHAH, N. 2012. Mental Agency and Metaethics. *In:* SHAFER-LANDAU, R. (ed.) *Oxford Studies in Metaethics.* Oxford: Oxford University Press.

EWING, A. C. 1948. *The Definition of Good,* London, Routledge and Kegan Paul.

FANTL, J. & MCGRATH, M. 2009. *Knowledge in an Uncertain World,* Oxford, Oxford University Press.

FELDMAN, R. 2000. The Ethics of Belief. *Philosophy and Phenomenological Research,* 60**,** 667-695.

FLETCHER, G. 2012. The Locative Analysis of Good For Formulated and Defended. *Journal of Ethics and Social Philosophy (JESP),* 6.

FOLEY, R. 1987. *The Theory of Epistemic Rationality,* Cambridge MA, Harvard University Press.

FOOT, P. 1978. *Virtues and Vices and Other Essays in Moral Philosophy,* Oxford, Oxford University Press.

FOOT, P. 2001. *Natural Goodness,* Oxford, Oxford University Press.

FRANKENA, W. K. 1973. *Ethics,* Englewood Cliffs, Prentice Hall.

FRICKER, M. 2009. The Value of Knowledge and The Test of Time. *Royal Institute of Philosophy Supplement,* 64**,** 121-138.

GEACH, P. T. 1956. Good and Evil. *Analysis,* 17**,** 33-42.

GETTIER, E. 1963. Is Justified True Belief Knowledge? *Analysis,* 23**,** 121-23.

GOLDMAN, A. 1976. Discrimination and perceptual knowledge. *Journal of Philosophy,* 73**,** 771-791.

GOLDMAN, A. & OLSSON, E. J. 2009. Reliablism and the Value of Knowledge. *In:* HADDOCK, A., MILLAR, A. & PRITCHARD, D. (eds.) *Epistemic Value.* Oxford: Oxford University Press.

GOLDMAN, A. I. 1999. *Knowledge in a Social World,* Oxford, Oxford University Press.

GOLDMAN, A. I. 2002. *Pathways to Knowledge: Private and Public,* Oxford, Oxford University Press.

GOLDMAN, A. I. 2015. Reliabilism, Veritism, and Epistemic Consequentialism. *Episteme,* 12**,** 131-143.

GRECO, J. 1999. Agent Reliabilism. *Philosophical Perspectives,* 13**,** 273-296.

GRECO, J. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity,* Cambridge, Cambridge University Press.

GRECO, J. 2011. The Value Problem. *In:* BERNECKER, S. & PRITCHARD, D. (eds.) *The Routledge Companion to Epistemology.* London: Routledge.

GRECO, J. 2012. A (Different) Virtue Epistemology. *Philosophy and Phenomenological Research,* 85**,** 1-16.

GRIMM, S. 2008. Epistemic Goals and Epistemic Values. *Philosophical and Phenomenological Research,* 77**,** 725-744.

GRIMM, S. 2009. Epistemic Normativity. *In:* HADDOCK, A., MILLAR, A. & PRITCHARD, D. (eds.) *Epistemic Value.* Oxford: Oxford University Press.

GRIMM, S. 2012. The Value of Understanding. *Philosophical Compass,* 7**,** 103-117.

HADDOCK, A., MILLAR, A. & PRITCHARD, D. (eds.) 2009. *Epistemic Value,* Oxford: Oxford University Press.

HORWICH, P. 2006. The Value of Truth. *Noûs,* 40**,** 347–360.

HURKA, T. 2001. *Virtue, Vice, and Value,* Oxford, Oxford University Press.

HYMAN, J. 2010. The Road to Larissa. *Ratio,* 23**,** 393-414.

JAMES, W. 1979. *The Will to Believe and Other Essays in Popular Philosophy,* Cambridge, MA, Harvard University Press.

JONES, W. E. 1997. Why Do We Value Knowledge? *American Philosophical Quarterly,* 34**,** 423-39.

KAGAN, S. 1998. Rethinking Intrinsic Value. *Journal of Ethics,* 2**,** 277-297.

KAPLAN, M. 1985. It's Not What You Know That Counts. *Journal of Philosophy,* 82**,** 350-363.

KELLY, T. 2003. Epistemic Rationality as Instrumental Rationality: A Critique. *Philosophy and Phenomenological Research,* 66**,** 612-640.

KELP, C. 2011. What's the Point of "Knowledge" Anyway? *Episteme,* 8**,** 53-66.

KELP, C. 2013. Extended Cognition and Robust Virtue Epistemology. *Erkenntnis,* 78**,** 245-252.

KORNBLITH, H. 2002. *Knowledge and its Place in Nature,* Oxford, Oxford University Press.

KORSGAARD, C. M. 1983. Two Distinctions in Goodness. *Philosophical Review,* 92**,** 169-195.

KRAUT, R. 2007. *What is Good and Why: The Ethics of Well-Being,* Cambridge, MA., Harvard university Press.

KVANVIG, J. 2008. Pointless truth. *Midwest Studies in Philosophy,* 32**,** 199-212.

KVANVIG, J. L. 2003. *The Value of Knowledge and the Pursuit of Understanding,* Cambridge, Cambridge University Press.

KVANVIG, J. L. 2014. Truth is Not the Primary Epistemic Goal. *In:* STEUP, M., TURRI, J. & SOSA, E. (eds.) *Contemporary Debates in Epistemology.* 2 ed. Oxford: Wiley-Blackwell.

LACKEY, J. 2007. Why We Don't Deserve Credit for Everything We Know. *Synthese,* 158**,** 345--361.

LACKEY, J. 2009. Knowledge and Credit. *Philosophical Studies,* 142**,** 27 - 42.

LEMOS, N. 1994. *Intrinsic Value: Concept and Warrant,* Cambridge, Cambridge University Press.

LITTLEJOHN, C. 2014. Introduction. *In:* LITTLEJOHN, C. & TURRI, J. (eds.) *Epistemic Norms: New Essays on Action, Belief and Assertion.* Oxford: Oxford: Oxford University Press.

LYNCH, M. P. 2005. *True to Life: Why Truth Matters,* Cambridge, MA, The MIT Press.

LYNCH, M. P. 2009a. Truth, Value and Epistemic Expressivism. *Philosophy and Phenomenological Research,* 79**,** 76-97.

LYNCH, M. P. 2009b. The Values of Truth and the Truth of Values. *In:* HADDOCK, A., MILLAR, A. & PRITCHARD, D. (eds.) *Epistemic Value.* Oxford: Oxford University Press.

MCHUGH, C. 2011a. Belief and Aims. *Philosophical Studies,* 160**,** 425-339.

MCHUGH, C. 2011b. Epistemic Deontology and Voluntariness. *Erkenntnis,* 77**,** 65-94.

MCHUGH, C. 2012. The Truth Norm of Belief. *Pacific Philosophical Quarterly,* 93**,** 8-30.

MCHUGH, C. 2013. Normativism and Doxastic Deliberation. *Analytic Philosophy,* 54**,** 447-465.

MCHUGH, C. 2014. Fitting Belief. *Proceedings of the Aristotelian Society,* 114**,** 167-187.

MIRACCHI, L. 2015. Competence to know. *Philosophical Studies,* 172**,** 29-56.

MOORE, G. E. 1993a. The Conception of Intrinsic Value. *In:* BALDWIN, T. (ed.) *Principia Ethica.* Revised ed. Cambridge: Cambridge University Press.

MOORE, G. E. 1993b. *Principia Ethica,* Cambridge, Cambridge University Press.

PILLER, C. 2014. What is Goodness Good For? *In:* TIMMONS, M. (ed.) *Oxford Studies In Normative Ethics.* Oxford: Oxford University Press.

PLANTINGA, A. 1993. *Warrant and Proper Function,* Oxford, Oxford University Press.

PRITCHARD, D. 2007. Recent Work on Epistemic Value. *American Philosophical Quarterly,* 44**,** 85-110.

PRITCHARD, D. 2010. Knowledge and Understanding. *In:* PRITCHARD, D., MILLAR, A. & HADDOCK, A. (eds.) *The Nature and Value of Knowledge: Three Investigations.* Oxford: Oxford University Press.

PRITCHARD, D. 2011. What is the Swamping Problem? *In:* REISNER, A. & STEGLICH-PETERSON, A. (eds.) *Reasons for Belief.* Cambridge: Cambridge University Press.

PRITCHARD, D. 2014. Truth as the Fundamental Epistemic Good. *In:* MATHESON, J. & VITZ, R. (eds.) *The Ethics of Belief.* Oxford: Oxford University Press.

PRITCHARD, D. & TURRI, J. 2014. *The Value of Knowledge* [Online]. Available: http://plato.stanford.edu/archives/spr2014/entries/knowledge-value/.

RABINOWICZ, W. & RØNNOW-RASMUSSEN, T. 2000. A Distinction in Value: Intrinsic and For Its Own Sake. *Proceedings of the Aristotelian Society,* 100**,** 33–51.

RIDGE, M. 2013. Getting Lost on the Road to Larissa. *Noûs,* 47**,** 181-201.

RIGGS, W. 2009. Getting the Meno Requirement Right. *In:* HADDOCK, A., MILLAR, A. & PRITCHARD, D. (eds.) *Epistemic Value.* Oxford: Oxford University Press.

RIGGS, W. D. 2002. Reliability and the Value of Knowledge. *Philosophy and Phenomenological Research,* 64**,** 79-96.

RIGGS, W. D. 2007. The Value Turn in Epistemology. *In:* HENDRICKS, V. & PRITCHARD, D. (eds.) *New Waves in Epistemology.* London: Palgrave Macmillan.

RØNNOW-RASMUSSEN, T. & ZIMMERMAN, M. J. 2005. Introduction. *In:* RØNNOW-RASMUSSEN, T. & ZIMMERMAN, M. J. (eds.) *Recent Work on Intrinsic Value.* Dordrecht: Springer.

ROSATI, C. S. 2006. Personal good. *In:* HORGAN, T. & TIMMONS, M. (eds.) *Metaethics After Moore.* Oxford: Oxford University Press.

ROSATI, C. S. 2009. Relational Good and the Multiplicity Problem. *Philosophical Issues,* 19**,** 205-234.

ROSS, W. D. 1939. *Foundations of Ethics,* Oxford, Oxford University Press.

RYAN, S. 2003. Doxastic Compatibilism and the Ethics of Belief. *Philosophical Studies,* 114**,** 47-79.

SARTWELL, C. 1992. Why Knowledge is Merely True Belief. *Journal of Philosophy,* 89**,** 167-180.

SETIYA, K. 2013. Epistemic Agency: Some Doubts. *Philosophical Issues,* 23**,** 179-198.

SHAH, N. 2003. How Truth Governs Belief. *Philosophical Review,* 112**,** 447-482.

SHAH, N. 2008. How Action Governs Intention. *Philosophers' Imprint,* 8**,** 1-19.

SHAH, N. & VELLEMAN, J. D. 2005. Doxastic Deliberation. *Philosophical Review,* 114**,** 497-534.

SMITHIES, D. 2012. The Normative Role of Knowledge. *NOÛS,* 46**,** 265-288.

SOSA, E. 1991. *Knowledge in Perspective,* Cambridge, Cambridge University Press.

SOSA, E. 2003. The Place of Truth in Epistemology. *In:* ZAGZEBSKI, L. & DEPAUL, M. (eds.) *Intellectual Virtue: Perspectives From Ethics and Epistemology.* Oxford: Oxford University Press.

SOSA, E. 2007. *A Virtue Epistemology,* Oxford, Oxford University Press.

SOSA, E. 2010. How Competence Matters in Epistemology. *Philosophical Perspectives,* 24**,** 465-475.

SOSA, E. 2011. *Knowing Full Well,* Princeton, NJ, Princeton University Press.

SOSA, E. 2013. Epistemic Agency. *Journal of Philosophy,* 110**,** 585-605.

STEGLICH-PETERSEN, A. 2006. Voluntarism and Transparent Deliberation. *South African Journal of Philosophy,* 25**,** 173-178.

STEGLICH-PETERSEN, A. 2009. Weighing the Aim of Belief. *Philosophical Studies,* 145**,** 395-405.

SWINBURNE, R. 2001. *Epistemic Justification,* Oxford, Oxford University Press.

SYLVAN, K. 2012. Truth Monism Without Teleology. *Thought: A Journal of Philosophy,* 1**,** 161-169.

THOMSON, J. J. 2001. *Goodness and Advice,* Princeton, NJ, Princeton University Press.

THOMSON, J. J. 2008. *Normativity,* Chicago and La Salle, Open Court.

TURRI, J., BUCKWALTER, W. & BLOUW, P. 2014. Knowledge and Luck. *Psychonomic Bulletin & Review,* 22**,** 378-390.

VAESEN, K. 2011. Knowledge without Credit, exhibit 4: Extended Cognition. *Synthese,* 181**,** 515-529.

WEDGWOOD, R. 2002. The Aim of Belief. *Philosophical Perspectives,* 16**,** 267-96.

WHITING, D. 2012a. Does Belief Aim (Only) at the Truth? *Pacific Philosophical Quarterly,* 93**,** 279-300.

WHITING, D. 2012b. Epistemic Value and Achievement. *Ratio,* 25**,** 216-230.

WHITING, D. 2013. The Good and the True (or the Bad and the False). *Philosophy,* 88**,** 219-242.

WHITING, D. 2014. Reasons for Belief, Reasons for Action, the Aim of Belief, and the Aim of Action. *In:* LITTLEJOHN, C. & TURRI, J. (eds.) *Epistemic Norms.* Oxford: Oxford University Press.

WILLIAMS, B. 1972. *Morality: An Introduction to Ethics,* Cambridge, Cambridge University Press.

WILLIAMSON, T. 2000. *Knowledge and its Limits,* New York, Oxford University Press.

ZAGZEBSKI, L. 1994. The Inescapability of Gettier problems. *Philosophical Quarterly,* 44**,** 65-73.

ZAGZEBSKI, L. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge,* Cambridge, Cambridge University Press.

ZAGZEBSKI, L. 2003. The Search for the Source of Epistemic Good. *Metaphilosophy,* 34**,** 12-28.

ZAGZEBSKI, L. 2004. Epistemic Value and the Primacy of What We Care About. *Philosophical Papers,* 33**,** 353-377.

ZIMMERMAN, M. J. 1999. In Defense of the Concept of Intrinsic Value. *Canadian Journal of Philosophy,* 29**,** 389-409.

ZIMMERMAN, M. J. 2001. *The Nature of Intrinsic Value,* Lanham, Rowman and Littlefield.

ZIMMERMAN, M. J. 2009. Understanding what's good for us. *Ethical Theory and Moral Practice,* 12**,** 429 - 439.

ZIMMERMAN, M. J. 2014. *Intrinsic vs. Extrinsic Value* [Online]. Available: http://plato.stanford.edu/archives/spr2015/entries/value-intrinsic-extrinsic/.